



INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ : H01J 17/30	A2	(11) International Publication Number: WO 99/27556 (43) International Publication Date: 3 June 1999 (03.06.99)
(21) International Application Number: PCT/US98/24963 (22) International Filing Date: 20 November 1998 (20.11.98) (30) Priority Data: 60/066,898 20 November 1997 (20.11.97) US 60/109,095 19 November 1998 (19.11.98) US (63) Related by Continuation (CON) or Continuation-in-Part (CIP) to Earlier Applications US 60/066,898 (CIP) Filed on 20 November 1997 (20.11.97) US 60/109,095 (CIP) Filed on 19 November 1998 (19.11.98) (71) Applicant (for all designated States except US): XACCT TECHNOLOGIES, INC. [US/US]; Suite 105, 2855 Kifer Road, Santa Clara, CA 95051 (US). (72) Inventors; and (75) Inventors/Applicants (for US only): WAGNER, Eran [US/US]; Apartement 217, 20677 Forge Way, Cupertino, CA 95014 (US); SCHWEITZER, Limor [IL/IL]; Achimeir Street 12/42, Ramat aviv gimel (IL); GIVOLY, Tal [IL/IL]; Barkan Street 2/1, 44288 Kfar-Saba (IL).		(74) Agent: RICHARDSON, Kent, R.; Wilson Sonsini Goodrich & Rosati, 650 Page Mill Road, Palo Alto, CA 94304-1050 (US). (81) Designated States: AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CU, CZ, DE, DK, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, UA, UG, US, UZ, VN, YU, ZW, ARIPO patent (GH, GM, KE, LS, MW, SD, SZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG). Published <i>Without international search report and to be republished upon receipt of that report.</i>
(54) Title: NETWORK ACCOUNTING AND BILLING SYSTEM AND METHOD (57) Abstract <p>In some embodiments, network traffic information is captured at network information sources. These sources provide detailed information about the network communications transactions at a network device. Importantly, different types of sources can provide different types of information. Gatherer devices gather the detailed information from the various information source devices and convert the information into standardized information. The gatherer devices can correlate the gathered information with account information for network transaction accounting. Manager devices manage the gatherer devices and store the gathered standardized information. The manager devices eliminate duplicate network information that may exist in the standardized information. The manager devices also consolidate the information. Importantly, the information stored by the manager devices represents the consolidated, account correlated, network transaction information that can be used for billing or network accounting. The system thereby provides a distributed network accounting and billing system.</p>		

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BK	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CJ	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KR	Republic of Korea	PT	Portugal		
CN	China	KZ	Kazakhstan	RO	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LI	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		
EE	Estonia						

NETWORK ACCOUNTING AND BILLING SYSTEM AND METHOD**RELATED APPLICATIONS**

5 This application relates to, and incorporates by reference, the United States patent application having serial number 60/066,898, entitled "Network Accounting and Billing System," having inventors Limor Schweitzer and Eran Wagner, filed November 20, 1997, and which has a common assignee. This application also relates to the United States patent application having serial
10 number XX/XXX,XXX, entitled "Network Accounting and Billing System," having inventors Limor Schweitzer and Eran Wagner, filed November 19, 1998, and which has a common assignee.

COPYRIGHT NOTICE

15 A portion of the disclosure of this patent document contains materials that are subject to copyright protection. The copyright owner has no objection to the facsimile reproduction by anyone of the patent disclosure, as it appears in the Patent and Trademark Office patents, files or records, but otherwise reserves all copyright rights whatsoever.

BACKGROUND OF THE INVENTION

A. Field of the Invention

This invention relates to the field of computer networks. In particular,
the invention relates to accounting and billing for services in a computer
5 network.

B. Description of the Related Art

The low cost of Internet connectivity and a wide range of services are
driving and more people onto the Internet, which is driving the deployment of
TCP/IP networks. This process has led to a new market of client-server
10 applications that enables the user to interact with other users and computer
systems around the world. The use of these applications is consuming more and
more Intranet and Internet bandwidth.

New applications such as "voice over IP (Internet Protocol)" and
streaming audio and video require even more bandwidth and a different quality
15 of service than email, or other less real-time applications. Also, the type quality
of service can vary according to the needs of the user. For example, typically,
businesses do not tolerate unavailable network services as easily as consumers.
Internet Service Providers (ISPs) therefore would like to price their available
bandwidth according to a user's needs. For example, flat monthly pricing may
20 be the best billing model for consumers, but businesses may want to be billed
according to their used bandwidth at particular qualities of service.

As ISPs continue to differentiate themselves by providing additional services, enterprise information technology managers will face similar problems to account for the escalating Intranet operating costs.

Therefore, ISPs and enterprise information technology managers will
5 want to account for session logging, bandwidth usage, directory data and application session information from a variety of sources.

Due to the diversity of IP data sources (e.g., routers, hubs etc.), the need for effect tracking far exceeds the problems addressed by telephone companies. Telephone companies track information such as circuit usage so it can be
10 correlated with account information. For example, businesses may use leased lines, consumers may have "Friends and Family" plans, cellular phones have different roamer fees according to the location of the user, etc. Typically, the phone company captures all of the data and uses batch processing to aggregate the information into specific user accounts. For example, all the long distance
15 calls made during a billing period are typically correlated with the Friends and Family list for each phone account at the end of a billing period for that account. This requires a significant amount of computing power. However, this type of problem is significantly simpler than attempting to track and bill for every transaction in an IP network. Therefore, what is desired is a system that allows
20 for accounting and billing of transactions on IP based networks.

The problem is even more difficult in IP network traffic because the information sources can exist and many different levels of the OSI network

model, throughout heterogeneous networks. Potential sources of information include packet use from routers, firewall authentication logging, email data, ISP session logging, and application layer use information. Therefore, what is desired is a system and method that track IP network usage information across

5 multiple layers of the OSI network model.

SUMMARY OF THE INVENTION

A network accounting and billing system and method are described. In some embodiments, the system can access any network related information sources such as traffic statistics provided by routers and switching hubs as well as application server access logs. The information can be accumulated in a central database for creating auditing, accounting and billing reports. Alternatively, the information can be sent directly to other systems such as rating engines used in customer care and billing systems.

In one embodiment, network traffic information is captured at network information sources (examples of information sources include network devices). These sources provide detailed information about the network communications transactions at a network device. Importantly, different types of sources can provide different types of information. Gatherer devices gather the detailed information from the various information source devices and convert the information into standardized information. The gatherer devices can correlate the gathered information with account information for network transaction accounting. Manager devices manage the gatherer devices and store the gathered standardized information. The manager devices eliminate duplicate network information that may exist in the standardized information. The manager devices also consolidate the information. Importantly, the information stored by the manager devices represents the consolidated, account correlated,

network transaction information used for billing. In addition to account information, transaction information can be correlated to other information such as geography information (e.g., the location of an accessed server) and/or transaction routing information (as may be used in peering agreements between Internet Service Providers). The system thereby provides a distributed network accounting and billing system.

In some embodiments, the gatherer devices can access sources through proxy gateways, firewalls, and/or address translation barriers.

In some embodiments, the gatherer devices can correlate the information about a specific transaction with a particular account by accessing the transaction's source and/or destination information. The source and/or destination information is then correlated with account information from an account information database.

BRIEF DESCRIPTION OF THE FIGURES

The figures illustrate the invention by way of example. The invention is not meant to be limited to only those embodiments of shown in the Figures. The same reference in different figures indicates the same element is being used in those figures.

Figure 1 illustrates a system including one embodiment of the invention.

Figure 2 illustrates an example of the data distillation used in the system of Figure 1.

Figure 3 illustrates data enhancements used in the data distillation.

Figure 4A illustrates example field enhancements that can be included in the data enhancements.

Figure 4B illustrates the creation of an enhanced record.

Figure 5 illustrates an example record merge.

Figure 6 illustrates an example of an alternative embodiment of the system.

DETAILED DESCRIPTION

A. System Overview

One embodiment of the system includes a multi-source, multi-layer network usage metering and mediation solution that gives Network Service
5 Providers (NSPs), including Internet Service Providers (ISPs) and enterprise network(Intranet) operators, the information needed to set the right-price for IP (Internet Protocol) services. With the system, the providers can generate accurate usage-based billing and implement usage-based charge-back models. The system derives IP session and transaction information, collected in real
10 time, from a multitude of network elements. The system gathers, correlates, and transforms data from routers, switches, firewalls, authentication servers, LDAP, Web hosts, DNS, and other devices to create comprehensive usage and billing records.

The system transforms raw transaction data from network devices into
15 useful billing records though policy-based filtering, aggregation, and merging. The result is a set of detail records (DRs). In some embodiments, the detail records are XaCCT Detail Records (XDRs™) available from XaCCT Technologies. DRs are somewhat similar in concept to the telephony industry's Call Detail Records (CDRs). Thus, DRs can be easily integrated with existing
20 Customer Care and Billing (CCB) systems.

In addition to billing data, DRs enable NSPs to deploy new services based on documented usage trends, plan network resource provisioning, and audit service usage. The system provides a clear picture of user-level network service use by tracking a variety of metrics such as actual session Quality of Service (QoS), traffic routes, and end-user application transactions.

The system is based on a modular, distributed, highly scalable architecture capable of running on multiple platforms. Data collection and management is designed for efficiency to minimize impact on the network and system resources.

The system minimizes network impact by collecting and processing data close to its source. Modular architecture provides maximum configuration flexibility, and compatibility with multiple network information sources.

The system, or other embodiments, may have one or more of the following features.

Data collection can be from a wide range of network devices and services, spanning all layers of the network - from the physical layer to the application layer.

Real-time, policy-based filtering, aggregation, enhancement and merging creates accurate, detailed and comprehensive session detail records (DRs).

Real time correlation of data from various sources allows billing record enhancement.

Leverages existing investment through integration with any customer care & billing solution, reducing costs, minimizing risks and shortened time-to-market.

5 Non-intrusive operation eliminates any disruption of network elements or services.

Web-based user interface allows off-the-shelf browsers to access the system, on-demand, locally or remotely.

Carrier-class scalability allows expansion to fit an NSPs needs without costly reconfiguration.

10 Distributed filtering and aggregation eliminates system capacity bottlenecks.

Efficient, centralized system administration allows on-the-fly system reconfigurations and field upgrades.

15 Customized reporting with built-in report generation or an NSPs choice of off-the-shelf graphical reporting packages.

Comprehensive network security features allow secure communication between system components and multiple levels of restricted access.

B. System Details

20 The following describes the system 100 of Figure 1. The system 100 allows NSPs to account for and bill for IP network communications. The following paragraphs first list the elements of Figure 1, then describes those elements and then describes how the elements work together. Importantly, the

distributed data gathering, filtering and enhancements performed in the system 100 enables load distribution. Granular data can reside in the peripheries of the system 100, close to the information sources. This helps avoid reduce congestion in network bottlenecks but still allows the data to be accessible from a central location. In previous systems, all the network information flows to one location, making it very difficult to keep up with the massive record flows from the network devices and requiring huge databases.

The following lists the elements of Figure 1. Figure 1 includes a number of information source modules (ISMs) including an ISM 110, an ISM 120, an ISM 130, an ISM 136, an ISM 140, and an ISM 150. The system also includes a number of network devices, such as a proxy server 101, a DNS 102, a firewall 103, an LDAP 106, a CISCO NetFlow 104, and a RADIUS 105. The system also includes a number of gatherers, such as a gatherer 161, a gatherer 162, a gatherer 163, a gatherer 164, and a gatherer 165. The system of Figure 1 also includes a central event manager (CEM) 170 and a central database (repository) 175. The system also includes a user interface server 185 and a number of terminals or clients 180.

This paragraph describes how the elements of Figure 1 are coupled. The various network devices represent devices coupled to an IP network such as the Internet. The network devices perform various functions, such as the proxy server 101 providing proxy service for a number of clients. Each network device is coupled to a corresponding ISM. For example, the proxy server 101 is

coupled to the ISM 110. The DNS 102 is coupled to the ISM 120. The firewall 103 is coupled to the ISM 130. The ISM 136 is coupled to the LDAP 106. The ISM 140 is coupled to the CISCO NetFlow 104. The ISM 150 is coupled to the RADIUS 105. Each gatherer is associated with at least one ISM. Thus, the
5 gatherer 161 is associated with the ISM 110 and is therefore coupled to that ISM. The gatherer 162 is coupled to the ISM 120. The gatherer 163 is coupled to the ISM 130 and the ISM 136. The gatherer 164 is coupled to the ISM 140. The gatherer 165 is coupled to the ISM 150. The various gatherers are coupled to the CEM 170. The user interface server is coupled to the terminals 180 and
10 the CEM 170.

The following paragraphs describe each of the various elements of Figure 1.

Network Devices

The network devices represent any devices that could be included in a
15 network. (Throughout the description, a network device, unless specifically noted otherwise, also refers to an application server.) A network device represents a subset of information sources that can be used by the system 100. That is, the network devices are merely representative of the types of sources of information that could be accessed. Other devices such as on-line transaction
20 processing databases can be accessed in other embodiments of the invention. Typically, the network devices keep logging and statistical information about their activity. A network information source can be the log file of a mail server,

the logging facility of a firewall, a traffics statistics table available on a router and accessible through SNMP, a database entry accessible through the Internet, an authentication server's query interface, etc. The network devices represent the information sources accessed by the ISMs.

5 Each type of network device can be accessing using a different method or protocols. Some generate logs while others are accessible via SNMP, others have proprietary APIs or use other protocols.

ISMs

10 The ISMs act as an interface between the gatherers and the network devices enabling the gatherers to collect data from the network devices. Thus, the ISMs represent modular, abstract interfaces that are designed to be platform-neutral. The information source modules act as interfaces or "translators", sending IP usage data, in real time, from the network devices to the gatherers. Each ISM is designed for a specific type of network data source. (In other

15 embodiments, some ISM are generic in that they can extract information from multiple network devices). ISMs can be packaged separately, allowing NSPs to customize ISM configurations to meet the specific requirements of their network. For example, in the system of Figure 1, if the NSP did not have Cisco NetFlow devices, then the ISM 140 would not have to be included.

20 The ISMs can communicate with its corresponding network device using protocols and formats such as UDP/IP, TCP/IP, SNMP, telnet, file access, ODBC, native API, and others.

In some embodiments, the reliability of system 100 is enhanced through on-the-fly dynamic reconfiguration, allowing the NSP to add or remove modules without disrupting ongoing operations. In these embodiments, the CEM 170 can automatically update the ISMs.

5 The following ISMs are available in some embodiments of the invention.

- Categorizer - Classifies a session to a category according to user-defined Boolean expression.
- DNS (e.g. ISM 120) - Resolves host names and IP addresses.
- 10 • Generic Proxy Server (e.g., ISM 110) - Collects data from access logs in a common log format.
- Port / Protocol Resolution - Converts protocol/port information to account names and vice versa.
- CheckPoint FireWall-1 - Collects data from FireWall-1 accounting log and security log.
- 15 • Cisco IOS IP Accounting - Collects accounting data from a Cisco router using IOS IP accounting.
- Cisco NetFlow Switching - Collects session data from a Cisco router via NetFlow switching.
- 20 • NETRANET – Collects information from a standard network device.
- Netscape Proxy Server - Collects data from a Netscape Proxy Server.

- Microsoft Proxy Server - Collects data from a Microsoft Proxy Server.

ISMs can be synchronous, asynchronous or pipe.

The data from an asynchronous ISM is dynamic so that the
5 asynchronous ISM reacts to the information and relays it to the associated
gatherer without prompting from other information sources in the system 100.
If the firewall 103 were a CheckPoint FireWall-1, then the ISM 130 would be
an example of an asynchronous ISM. When a network session is initiated, the
details are recorded by the FireWall-I 103. The corresponding ISM 130 receives
10 the details and passes them on automatically to the gatherer 163.

Synchronous ISMs provide its information only when accessed by a
gatherer. The ISM 120 is an example of a synchronous ISM. The DNS server
102 maintains information matching the IP addresses of host computers to their
domain addresses. The ISM 120 accesses the DNS server 102 only when the
15 ISM 120 receives a request from the gather 162. When the DNS server 102
returns a reply, the ISM 120 relays the reply information to the gatherer 162.

Pipe ISMs operate on record flows (batches of records received from
information sources). Pipe ISMs process one or more enhancement flows the
records as the flows arrive. The pipe ISM may initiate new record flows or may
20 do other things such as generate alerts or provision network elements to provide
or stop services. The pipe is implemented as an ISM to keep the internal
coherency and logic of the architecture. (Record flows can terminate in a

database or in a pipe ISM. The pipe ISM can perform filtering and aggregation, send alarms, or act as a mediation system to provision network elements when some event occurs or some accumulated value is surpassed. Specifically, pipe ISMs can act to enable pre-payment systems to disable certain services such as a voice IP call, when the time limit is surpassed or amount of data is reached.)

The gatherers can include caches and buffers for storing information from the ISMs. The buffers allow the gatherers to compensate for situations where there is a loss of connection with the rest of the system 100. The cache sizes can be remotely configured. The cache minimizes the number of accesses to the Information Source.

ISM queries can be cached and parallelized. Caching of synchronous ISM queries provides for fast responses. Parallelizing queries allows for multiple queries to be processed at the same time.

Gatherers

The gatherers gather the information from the ISMs. In some embodiments, the gatherers are multi-threaded, lightweight, smart agents that run on non-dedicated hosts, as a normal user application on Windows NT or Unix, as a background process, or daemon. What is important though is that the gatherers can be any hardware and/or software that perform the functions of a gatherer.

The gatherers can be installed on the same network segment as the network device such as router and switch or on the application server itself. This placement of a gatherer minimizes the data traffic impact on the network.

The gatherers collect network session data from one or more ISMs.

5 Session data can be sent to another gatherer for enhancement or to the CEM 170 for merging and storing in the central database 170. The gatherers can be deployed on an as needed basis for optimal scalability and flexibility.

The gatherers perform flexible, policy-based data aggregation. Importantly, the various types of ISMs provide different data and in different
10 formats. The gatherers normalize the data by extracting the fields needed by the CEM 170 and filling in any fields that may be missing. Thus, the gatherers act as a distributed filtering and aggregation system. The distributed data filtering and aggregation eliminates capacity bottlenecks improving the scalability and efficiency of the system 100 by reducing the volume of data sent on the network
15 to the CEM 170.

Aggregation can be done by accumulating groups of data record flows, generating a single data record for each group. That single record then includes the aggregated information. This reduces the flow of the data records.

Filtering means discarding any record that belongs to a group of
20 unneeded data records. Data records are unneeded if they are known to be collected elsewhere. A policy framework enables the NSP to configure what to collect where.

Filtering and/or aggregation can be done at any point along a data enhancement (described below) so that aggregation schemes can be based on enhanced data records as they are accumulated. The filtering and/or aggregation points are treated by the system 100 as pipe ISMs which are flow termination and flow starting points (ie: like an asynchronous ISM on the starting end and like a database on the terminating end). Data enhancement paths and filtering and/or aggregation schemes can be based on accumulated parameters such as user identification information and a user's contract type.

As noted above, the PISM can be used in the context of filtering and/or aggregation. One or more record flows can terminate at the PISM and can be converted into one or more new record flows. Record flows are grouped based on matching rules that apply to some of the fields in the record flows, while others are accumulated or undergo some other operation such as "maximum" or "average". Once the groups of accumulated records have reached some threshold, new accumulated records are output. This can be used for example in order to achieve a business-hybrid filtering and aggregation data reduction by imposing the business rules or the usage-based products that are offered to the customer, onto the record flows as they are collected in real-time. This is done instead of previous system where the information is stored in a database and then database operations are performed in order to create bills or reports. The filtering and aggregation reduces the amount of data that is stored in the central

database 175 while not jeopardizing the granularity of data that is necessary in order to create creative usage-based products.

Typically, data collected from a single source does not contain all the information needed for billing and accounting, such as user name and organization. In such cases, the data is enhanced. By combining IP session data from multiple sources, such as authentication servers, DHCP and Domain Name servers, the gatherers create meaningful session records tailored to the NSP's specific requirements. In the example of Figure 1, the gatherer 161 can provide information to the gatherer 162 so that the source IP address for an Internet session from the proxy server 101 can be combined with the domain address from the DNS server 102.

The enhancement procedure can be triggered by an asynchronous ISM. The information from the asynchronous ISM is associated with field enhancements in the central database 175. A field enhancement defines how a field in the central database is filled from the source data obtained from the asynchronous ISM. Through the field enhancements, the missing parameters are added to a record using the data collected from one or more synchronous ISMs. Enhancements are described in detail below.

The gatherers can include caches and buffers for storing information from the ISMs. The buffers allow the gatherers to compensate for situations where there is a loss of connection with the rest of the system 100. The caches

can reduce the number of accesses to an information source. The buffer and/or cache sizes can be remotely configured.

Central Event Manager (CEM)

5 The Central Event Manager (CEM) 170 acts as the central nervous system of the system 100, providing centralized, efficient management and controls of the gatherers and the ISMs.

The CEM 170 can perform one or more of the following tasks:

- Coordinates, controls, and manages the data collection process. The CEM 170 coordinates the operation of the gatherers and manages the flow of data through the system 100 through the collection scheme defined in the system configuration. The latter includes the configuration of the gatherers, the ISMs, the network devices, the fields in the central database 175 (described below), and the enhancement procedures. Based on the collection scheme the CEM 170 determines the system 100's *computation flow* (the set of operations the system 100 must perform to obtain the desired information). The CEM 170 then controls all the gatherers, instructing them to perform, in a particular sequence, the operations defined in the computation flow. The CEM 170 receives the records collected by the gatherers and stores them in the central database 175. NSPs can configure the CEM 170 to *merge* duplicate records

before storing them in the central database 175. Record merging is described below.

- Performs clean-up and aging procedures in the database 175. The system 100 collects and stores large amounts of session information every day. The CEM 170 removes old data to free space for new data periodically. The NSP defines the expiration period for the removal of old records. The CEM 170 is responsible for coordinating the removal of records from the central database 175. The CEM 170 places a time stamp on every record when the record enters the central database 175 and deletes the record after the time period the NSP has defined elapses.
- Provides centralized system-wide upgrade, licensing, and data security. The NSP can perform version upgrades of the system 100 at the CEM 170. The gatherers can be automatically upgraded once a new version is installed on the host computer of the CEM 170. ISMs are also installed via the CEM 170 and exported to the gatherers. The CEM 170 maintains a list of licenses installed in the system and verifies periodically if the system is properly licensed. This feature lets the NSP centrally install and uninstall licenses. It also prevents unlicensed use of the system 100 and any of its components.
- Monitors the state of the gatherers and ISMs. The gatherers periodically communicate with the CEM 170. The CEM 170

continuously monitors the state of each gatherer and network devices in the system 100. The CEM 170 can be fault-tolerant, that is, it can recover from any system crash. It coordinates the recovery of the system 100 to its previous state.

5 In some embodiments, a key directory server is associated with the CEM 170. To transfer less data between the elements of the system 100, it is desirable that each piece of data to carry little descriptive data. For example, if IP address data is transferred between a gatherer and the CEM 170, a description of the IP address data is typically included. In some embodiments, data name/key, type, and length descriptions are included with the actual IP address data. In other
10 embodiments, there the key directory server reduces the amount of descriptive information being sent. Every key in the directory server has a type and a length. Fields can be identified as variable length. Therefore, data type information need not be transmitted between elements in the system 100 if the
15 elements use a common reference key stored in the directory server. Returning to the IP address data, by using the key directory server, elements need only send two bytes for the key id and four bytes for the actual address. Most of the data being sent in the system is relatively short in length. Therefore, the directory server helps reduce the amount of information being sent between the
20 elements in the system 100.

Keys can be added to the directory server. The directory server can therefore support expansion of the kinds of fields being sent by allowing system

elements to update their locally stored key ids. For example, after a recipient receives a record with an "unknown" key, it contacts the directory server to get the key definition.

Central Database

5 The central database 175 is the optional central repository of the information collected by the system 100. The central database 175 is but one example of a sink for the data generated in the system 100. Other embodiments include other configurations. The central database 175 stores and maintains the data collected by the gatherers, as well as the information on the configuration
10 of the system 100. Thus, in configuring the system 100, the NSP defines what data will be stored in each field in the central database 175 and how that data is collected from the ISMs.

 The information on network sessions is stored in the database in the form of a table. Each field in the table represents a network session parameter.
15 Each record describes a network session. The system 100 has a set of predefined fields that are configured by the CEM 170 on installation. The NSP can modify the central database 175 structure by adding, deleting, or modifying fields. The NSP access the data in the central database 175 by running queries and reports. The old data is removed from the central database 175 to free space
20 for new data periodically. You can specify the time interval for which records are stored in the central database 175. The structure of the central database 175 with some of the predefined fields is illustrated in the following figure.

As each IP session may generate multiple transaction records, during the merge process the CEM 170 identifies and discards duplications, enhancing the efficiency of the data repository. Generally, data records are passed through the merger program, in the CEM 170, into the central database 175. However, the data records are also cached so that if matching records appear at some point, the already stored records can be replaced or enhanced with the new records. The database tables that contain the record flows can be indexed, enhancing the efficiency of the data repository. A merge is achieved by matching some of the fields in a data record and then merging the matching records from at least two record flows, transforming them into one record before updating the central database 175. In some embodiments, adaptive tolerance is used to match records. Adaptive tolerance allows for a variation in the values of fields that are compared (e.g., the time field value may be allowed to differ by some amount, but still be considered a match). The adaptive aspect of the matching can include learning the appropriate period to allow for the tolerance. The reason that the records that do not match any previous records are sent through into the central database 175, in addition to being cached for later matching, is to avoid loss of data in case of system failure.

The following table illustrates an example of the types of records stored in the central database 175 by the CEM 170.

Source IP	Destination IP	Source Host	Destination Host	Service	Date/Time	Duration	Total Bytes	Counter
199.203.13.2.187	204.71.177.35	pcLev.xaccr.com	yahoo.com	http	1998-04-26 10:56:55	6464	435666	261019
199.203.13	207.68.137.5	prodigy.xac	microsoft.co	telnet	1998-04-26	747	66743	261020

2.131	9	et.com	m		10:56:55			
199.203.13	199.203.132	pcEitan.xac	xpert.com	smtp	1998-04-26	82	33667	261021
2.177	1	et.com			10:56:55			
199.203.13	204.162.80.1	pcAdi.xacc	cnet.com	http	1998-04-26	93	33567	261022
2.173	82	t.com			10:56:55			

The system 100 supports a non-proprietary database format enabling the central database 175 to run on any of a number of commercially available databases (e.g., MS-SQL Server, Oracle Server, DB2, etc.).

5 User Interface Server and Clients

The User Interface Server (UIS) 185 allows multiple clients (e.g. terminals 180) to access the system 100 through, the Microsoft Internet Explorer with Java™ Plug-in or Netscape Navigator with Java™ Plug-in. Other embodiments can use other applications to access the system 100. The main

10 function of the UIS 185 is to provide remote and local platform independent control for the system 100. The UIS 185 can provide these functions through windows that correspond to the various components of the system 100. Access to the system 100 can be password protected, allowing only authorized users to log in to the system and protecting sensitive information.

15 The NSP can perform one or more of the following main tasks through the UIS 185:

- Configure the system 100.
- Create and run queries and reports on network activity and resource consumption.

- Register and license the system 100.

C. Data Distillation

Figure 2 illustrates the data distillation process performed by the system of Figure 1. The data distillation aggregates and correlate information from many different network devices to compile data useful in billing and network accounting.

First, the ISMs 210 gather data from their corresponding network device. Note that for some ISMs (e.g. pipe ISMs), real-time, policy-based filtering and aggregation 215 can also be done. This data is then fed to the gatherers 220. The gatherers 220 perform data enhancement to complete the data from the ISMs 210. The results are provided to the CEM 170. The CEM 170 performs data merges 270 to remove redundant data. The merged data is then optionally stored in the central database 175 as a billing record 275 or is sent directly to an external system. The billing record information can be accessed from external applications, through the application interface 290, via a data record 280. Filtering and/aggregation and/or data enhancements can be done at any stage in the system 100.

D. Data Enhancement

As mentioned above, the gatherers 220 provide data enhancement features to complete information received from the ISMs 210. The following

describes some example data enhancement techniques used in some embodiments of the invention.

Figure 3 illustrates an example of data enhancement. Data enhancement comprises a number of field enhancements. A field enhancement specifies how
5 the data obtained from the trigger of the enhancement procedure is processed before it is placed in a single field in the central database 175. The data can be placed in the field directly, or new information may be added to the record by applying a Synchronous ISM function. (In the example below, the function is "resolve the IP address to a host FQDN"). Field enhancements may involve one
10 or multiple steps. There is no limit to the number of steps in a Field Enhancement. The data record starts with fields obtained from an asynchronous ISM 300. The fields in the DR 300 are then enhanced using the field enhancements. The enhanced fields result in the DR 320.

A visual representation of an enhancement can be presented to the NSP.
15 The enhancement may include an itinerary of ISMs starting off with an AISM, passing through PISMs, and terminating in the CEM 170. Using this view of the system 100, the NSP need not be shown the actual flow of data since the flow may be optimized later in order to achieve better performance. This is more of a graphical logical view of how the enhancement is achieved in steps. (PISMs can
20 terminate more than one flow and initiate more than one flow.)

A visual representation of a field enhancement shows the per-field flow of data correlation. This process ends in the CEM 170 or in a PISM. The NSP

supplies information telling the system 100 how to reach each of the terminating fields (in the CEM 170 or the PISM) starting off from the initiating fields (PISM or AISM). Each step of enhancement defines cross correlation with some SISM function.

5 Figure 4A illustrates various field enhancements (410 through 440). A field enhancement includes applying zero or more functions to a field before storing the field in a specified field in the central database 175.

One-step Field Enhancement 410. The initial source data from the asynchronous ISM is placed directly in a field in the central database 175.

10 Example: the field enhancement for the Source IP field.

Two-step Field Enhancement 420. The initial source data from the asynchronous ISM is used to obtain new additional data from a synchronous network device and the new data is placed in a field in the central database 175.

Example: the field enhancement for the Source Host field.

15 Three-step Enhancement 430. The initial source data from the asynchronous ISM is used to obtain additional data from a synchronous ISM. The result is used to obtain more data from another ISM and the result is placed in a field in the central database 175.

20 The following illustrates an example data enhancement. Suppose the data obtained from a proxy server 101 contains the source IP address of a given session, such as 199.203.132.2, but not the complete domain address of the host computer (its Fully Qualified Domain Name), such as www.xacct.com. The

name of the host can be obtained by another network device - the Domain Name System (DNS 102) server. The DNS server 102 contains information that matches IP addresses of host computers to their Fully Qualified Domain Names (FQDNs). Through an enhancement procedure the information collected from the proxy server 101 can be supplemented by the information from the DNS 102. Therefore, the name of the host is added to the data (the data record) collected from the proxy server 101. The process of adding new data to the data record from different network devices can be repeated several times until all required data is collected and the data record is placed in the central database 175.

Figure 4B illustrates another example data enhancement where an enhanced record 490 is created from an initial netflow record 492. Fields in the enhanced record 490 are enhanced from the radius record 494, the QoS policy server record 496, the NMS DB record 498, and the LDAP record 499.

15 Defining Enhancement Procedures

The following describes the process for defining enhancement procedures in some embodiments of the system. Typically defining an enhancement procedures for the system 100 includes (1) defining enhancement procedures for each asynchronous ISM and (2) configuring field enhancements for all fields in the central database 175 for which the NSP wants to collect data originating from an asynchronous ISM that triggers the corresponding enhancement procedure.

An enhancement procedure can be defined as follows:

1. Access the CEM 170 using the UIS 180.
2. Select the enhancement procedures list using the UIS 180.
3. Define the name of the new enhancement procedure.
- 5 4. Select a trigger for the new enhancement procedure. The trigger can correspond to any asynchronous ISM in the system 100.

Alternatively, the trigger can correspond to any asynchronous ISM in the system 100 that has not already been assigned to an enhancement procedure.
- 10 5. Optionally, a description for the enhancement procedure can be provided.
6. The new enhancement procedure can then be automatically populated with the existing fields in the central database 175.

Optionally, the NSP can define the fields (which could then be
15 propagated to the central database 175). Alternatively, based upon the type of asynchronous ISM, a preset set of fields could be proposed to the NSP for editing. What is important is that the NSP can define field procedures to enhance the data being put into the data records of the central database 175.
- 20 7. The NSP can then define the field enhancements for every field in the new enhancement procedure for which the NSP wants to collect

data from the ISM that is the trigger of the new enhancement procedure.

Defining Field Enhancements

Defining a field enhancement involves specifying the set of rules used to fill a database field from the information obtained from the trigger of the enhancement procedure. The NSP defines field enhancements for each field in which NSP wants to collect data from the trigger. If no field enhancements are defined, no data from the trigger will be collected in the fields. For example, suppose the firewall asynchronous ISM 130 that triggers an enhancement procedure. Suppose the central database 175 has the following fields: source IP, source host, destination IP, destination host, user name, total bytes, service, date/time, and URL. If the NSP wants to collect session data for each field except the URL from the firewall ISM 130, which triggers the enhancement procedure, the NSP defines a field enhancement for each field with the exception of the URL.

In some embodiments, the field enhancements are part of the enhancement procedure and the NSP can only define and modify them when the enhancement procedure is not enabled.

The field enhancements can be defined in a field enhancement configuration dialog box. The field enhancement configuration dialog box can have two panes. The first displays the name of the enhancement procedure, the name of its trigger, and the name and data type of the field for which the NSP is

defining the field enhancement. The second is dynamic and interactive. Its content changes depending on the NSP's input. When first displayed, it has two toggle buttons, End and Continue, and a list next to them. The content of the list depends on the button depressed.

5 When End is depressed, the list contains all output fields whose data type matches the data type of the field for which the NSP is defining the field enhancement. For example, if the field's data type is IP Address, the list contains all fields that are of the same type, such as source IP and destination IP that the AISM supplies. The fields in the list can come from two sources: (1) the
10 source data which the gatherer receives from the trigger and (2) the result obtained by applying a synchronous ISM function as a preceding step in the field enhancement. The following notation is used for the fields:

OutputFieldName for the output of a field origination from the trigger

SISName.FunctionName(InputArgument).OutputField for the output of a

15 field that is the result of applying a function

SISName...OutputField for the output of a field that is the result of

applying a function as the final step of a field enhancement

The following examples are presented.

Source IP is the field provided by the trigger of the enhancement

20 procedure that contains the IP address of the source host.

DNS...Host Name and DNS.Name(Source IP).Host name are the names

of a field originating from the resolved function Name of a network device

called DNS that resolves the IP address to a domain address. The input
argument of the function is the field provided by the trigger of the enhancement
procedure, called source IP. It contains the IP address of the source host. The
function returns the output field called Host Name that contains the domain
5 address of the source host. The notation DNS...Host Name is used when the
field is the result of applying the function as the final step of a field
enhancement. The notation is DNS.Name(Source IP).Host Name is used when
the field is used as the input to another function.

10 In the user interface, if End is unavailable, none of the output fields
matches the data type of the field.

When Continue is depressed, the list contains all applicable functions of
the available synchronous network device configured in the system 100. If the
preceding output does not match the input to a function, it cannot be applied and
does not appear on the list.

15 The following notation is used for the functions:

SISName.FunctionName(InputFieldName:InputFieldDataType)
→ (OutputFieldName:OutputFieldDataType)

When the function has multiple input and/or output arguments, the
notation reflects this. The arguments are separated by commas.

20 The following example shows a field enhancement.

DNS. Address(Host Name:String) → (IP Address:IP Address)

Where DNS is the name of the synchronous ISM (or network device) as it appears in the system configuration.

Address is the name of the function.

(Host Name:String) is the input to the function - host FQDN of data type

5 String

(IP Address:IP Address) is the output - IP address of data type IP

Address

The NSP can define the field enhancement by choosing items from the list. The list contains the option <none> when the End button is depressed.

10 Choosing this option has the same effect as not defining a field enhancement: no data from the trigger will be stored in the field in the central database 175.

At your request, here is one paragraph describing the algorithm(s) transforming the user representation of an Enhancement Procedure to an actual Computation Flow performed by the system. The source code accompanying
15 this actually implements most of this. Of course, that code relies on the rest of the software (lots of other code), that may be required to completely understand it.

Field Enhancement Computation

In some embodiments, an enhancement procedure (EP) defines a flow of
20 data records in the system 100, originating from an AISM that serves as a trigger, and terminating at the a database table, or a PISM (Pipe ISM - has the behavior of the database of being assigned input keys in much the same way,

but being installable/replacable logic in the system after it has been deployed), known as the Target of the EP. An EP is represented to the user as a collection of Field Enhancements (FEs). Each FE identifies what is performed in order to obtain a value (or not, if it is not defined), for that Field (either database field, or
5 input key to PISM), using values originating from the Trigger AISM, possibly, utilizing functions of SISMs available in the system. The NSP defines every FE independently of other FEs, and completely defines what value is placed in the defined field.

In contrast to the above description of the presentation to the user of an
10 EP, the system 100 (e.g., the CEM 170) computes what is called a Computation Flow, which is the actual operational instructions and parameters that are instructed to the gatherer's to achieve the desired NSP result (as defined in the EP). The Computation Flow is the detailed instructions required to achieve a "snowball" affect, of having data originating from a Trigger (AISM) in one
15 gatherer , undergo specific enhancements by adding results from invocations of functions of SISMs with specific input parameters on the same gatherer or other gatherers (as required), and removing unused fields along this same path, eventually storing the user-configured results in the user-configured fields in an optimized fashion. The optimizations applied are:

- 20 • Transfer minimal data between any elements in the system;

- Do not compute the same function on the same input more than once. In other embodiments, this optimization may be changed to allow for sharing of intermediary values of field enhancements;
- Perform local enhancements before going on to a remote gatherer;
- 5 • Apply a "cost" function to transferring along communication links and traverse the minimal cost in a computation flow required to achieve correct results.

This set of optimizations can be performed, in some embodiments, using the code in Appendix A. That is, the transformation (from EP to Computation
10 Flow, or simply Flow) is achieved by coupling together eleven algorithms and data transformations one after the other. These are described in the code module attached.

Note: 1) Within the code, EP, Rule, Computation Flow are used in some cases to refer to the same objects. 2) The eleven transformations that are applied
15 to the user representation of the EP, after being stored in a database structure, are outlined beginning at line 124 of the attached source module "FlowManager.java".

E. Record Merges

Figure 5 illustrates an example record merge. Record merging removes
20 duplicate records from the central database 175.

The following example shows how merges work and illustrates the need for merging duplicate records. Suppose the system 100 is using two

asynchronous ISMs 110 and 130. All outbound network traffic going through the proxy server 101 is routed through the firewall 103. The firewall 103 records the proxy server 101 as the source of all sessions passing through the proxy server 101, although they originate from different workstations on the network.

5 At the same time, the proxy server 101 records the destination of all sessions as the firewall 103, although their actual destinations are the different Internet sites.

Therefore, all sessions are logged twice by the system 100 and the records are skewed. The data from the firewall 103 indicates the destination of a given session, but not the source (see data record 520), while the data from the proxy server 101 records the source, but not the destination (see data record 510). Defining a merge eliminates the duplication of records.

A merge can be defined instructing the CEM 170 to store the destination data obtained from the firewall 103 and the source data from the proxy server 101 in the central database 175. The merge will also eliminate the problem of skewed data by storing the correct source and destination of the session in the central database 175. Both network devices provide information on the URL. The latter can be used to identify the fact that the two seemingly independent records (510 and 520) are actually two logs of the same session.

20 Two enhancement procedures are defined for the example of Figure 5. The trigger of the first, designated Flow One, is the Proxy Server Asynchronous Information Source Module. The trigger of the second, Flow Two, is the

Firewall Asynchronous Information Source Module. The records from Flow One and Flow Two are records of the same session. They both have the same value for the URL field. Based on this value, the CEM 170 identifies the two records are double logs of the same session. It merges the two data records taking the Source IP value from Flow One and the Destination IP from Flow Two as the values to be stored in the central database 175.

Defining Merges

The following describes defining merges. A merge is a set of rules that specify how duplicate records from multiple enhancement procedures must be identified and combined before being stored in the central database 175. The NSP can merge the records from two or more enhancement procedures. To define a merge, the NSP identifies the following information.

- The enhancement procedures included in the merge.
- How to identify duplicate records (which fields of the records must match).
- How to combine the records; that is, for each field, which value (from which enhancement procedure) must be stored in the central database 175. (Optional)

If the NSP does not specify how records must be combined, the records are merged as follows:

- When the values in all but one of the fields are null, the non-null value is stored.

- When the fields contain non-null values, the value of the first record received (chronologically) is stored.

F. Additional Embodiments

The following describes additional embodiments of the invention.

5 In some embodiments, the user interface used by an NSP to configure the system 100 can be presented as a graphical representation of the data enhancement process. Every step in the enhancement can be shown as a block joined to another block (or icon or some graphical representation). The properties of a block define the operations within the block. In some
10 embodiments, the entire data enhancement process from network devices to the central database 175 can be shown by linked graphics where the properties of a graphic are the properties of the enhancement at that stage.

In some embodiments, multiple CEMs 170 and/or central databases 175 can be used as data sources (back ends) for datamart or other databases or
15 applications (e.g., customer care and billing systems).

In some embodiments, the types of databases used are not necessarily relational. Object databases or other databases can be used.

In some embodiments, other platforms are used. Although the above description of the system 100 has been IP network focussed with Unix or
20 Windows NT systems supporting the elements, other networks (non-IP networks) and computer platforms can be used. What is important is that some

sort of processing and storing capability is available at the gatherers, the CEMs, the databases, and the user interface servers.

In some embodiments, the gatherers and other elements of the system 100, can be remotely configured, while in other embodiments, some of the
5 elements need to be configured directly. For example, a gatherer may not be remotely configurable, in which case, the NSP must interface directly with the computer running the gatherer.

In other embodiments, the general ideas described herein can be applied to other distributed data enhancement problems. For example, some
10 embodiments of the invention could be used to perform data source extraction and data preparation for data warehousing applications. The gatherers would interface with ISMs that are designed to extract data from databases (or other data sources). The gatherers would perform filtering and aggregation depending upon the needs of the datamart (in such an embodiment, the central database and
15 CEM could be replaced with/used with a datamart). The data enhancement would then be done before storing the information in the datamart.

Figure 6 illustrates a system 600 where multiple systems 100 are linked together. This system could be an ISPs point of presence accounting system. The system 620 and the system 610 can store detailed network accounting
20 information in their local detailed accounting databases. This information can then be aggregated and sent over the more expensive long distance links to the billing database in the system 630. Customer service information can still be

accessed at the detailed accounting database, but the aggregated information may be all that is needed to create the bills.

G. Conclusions

A network accounting and billing system and method has been
5 described. In some embodiments, the system can access any network related
information sources such as traffic statistics provided by routers and switching
hubs as well as application server access logs. These are accumulated in a
central database for creating auditing, accounting and billing reports. Because of
the distributed architecture, filtering and enhancements, the system efficiently
10 and accurately collects the network usage information for storage in a form that
is useful for billing and accounting.

Appendix A - FlowManager

Example

```

5 //*****
//
// NAME & TITLE:
// FlowManager Analyzes the various Enhancement Rules and creates Computation
// Flows
//
10 // ABSTRACT:
// - Analyzes the various Enhancement Rules and creates Computation Flows
// - Respond to Gatherer Requests for Instructions
// - Instructions are stored in UNIRInstruction, which can either be an element
// of a
15 // vector or a vector of UNIRInstructions in itself
// - Allow for ReComputation of Particular Flows
// - Implement Helper methods for finding interdependencies between Enhancement
// Rules
// and ISs that are referenced by them. Important when attempting to delete
20 an IS
// or when an IS configuration has changed, to see which Enhancements may be
// affected.
//
// this class is closely related to the Gatherer.EnhanceCompFlowStep class,
25 which acts
// on the instructions created here.
// An instruction needs to know the following:
// 1. What type it is (receive, enhance, send or other)
// 2. If it is of type enhance:
30 // 2.1 Where are the Input Units for the Enhancement Function
// ("where" means position in the Arguments)
// 2.2 What IS and what function to perform
// 2.3 Where in the Arguments to put the function results.
// 3. If it is of type send, needs to know to whom to send. CEM or another
35 Gatherer.
//
// For further understanding, see Gatherer.EnhanceCompFlowStep
//
// CHANGE HISTORY:
40 //
// Programmer Change Date Change Description
// -----
// Adi Gavish 61/12/97 Initial version
//
45 //*****
import java.util.*;

```

```

import Util.*;
import Util.Persistence.*;
import Util.UNIRTypes.*;
import Util.Structs.*;
5  import Util.Log.*;
import COM.objectspace.jgl.*;

//*****
public class FlowManager
10  {
    // Current Rule we are processing
    int m_RuleID ;

    // Pipe ISM variables
15  boolean m_IsTargetPipe = false ;
    int m_TargetIS      = 0 ;
    int m_TargetNode     = 1 ;
    int m_TargetGatherer = 0;

20  // Holds all Enhancements for a Rule. All elements are of type
    EnhanceOpNode.
    public EnhanceOpSList m_OpList = new EnhanceOpSList();
    public EnhanceOpSList m_OpTmpList = new EnhanceOpSList();

25  // All output parameters are added to this list, with their positions in
    the list
    // used to define the actual position of UNIRs in the Data Arguments
    collected by
    // the ISs of the flow.
30  // After Output positions are defined, the input parameters are updated
    with
    // the new positions (So the Enhancement step in the Gatherer will know
    where to
    // get the input parameters of some enhancement function)
35  // see allocateUNIROutputData method
    SList UNIRAllocation ;

    // Temporary Storage of instructions converted from Enhancements
    SList m_TmpInstructions ;

40  // HashTable For Flow Instruction Storage, access by RuleID Stores a
    complete set
    // of instructions for every Flow (== Every Rule)
    Hashtable m_Instructions = new Hashtable();

45  // Offset to fix UNIR Indexes (Take FlowID, Step into account)
    // See Placeholders in Gatherer.EnhanceCompFlowStep
    final static int FLOW_OFFSET = 2;

50  // Global counter used in DFS Topology Sort

```

```

        int m_TimeCtr = 0 ;

        //*****
        // Build All Flows
5       // ~ Loop on all rules
        // ~ build an instruction vector for each Rule
        // ~ Store set of instructions per Rule in m_Instructions
        //*****

        public void computeFlows()
10       {
            Rules RuleList = OEM.instance().getRules();
            IF (RuleList != null)
            {
                // Loop on all Rules
15         for(Enumeration El = RuleList.elements(); El.hasMoreElements(); )
                {
                    // Compute Flow for specific Rule (If Enhancement is
                    Enabled:
                    Rule CurrentRule = ( (Rule){El.nextElement()} );
20         if (( CurrentRule != null ) && (CurrentRule.getEnabled()
                    == true: )
                    {
                        int RuleID = CurrentRule.getRuleID();
                        computeRuleFlows(RuleID);
25         }
                }
            }
        }

30     //*****
        // Compute Flow for RuleID. At the end of this method, All
        FlowInstructions for RuleID will be
        // generated and saved as an SList in a Hash Table, accessed by RuleID
        //
35     // parameters
        // RuleID - ID of rule to build the instruction set for
        //*****
        public void computeRuleFlows(int RuleID)
        {
40         Log.instance().writeLog(new LogDebug("Compute Flow for Rule "+RuleID));
            m_RuleID = RuleID ;

            try
            {
45         // 1. Build Initial EnhanceOpNode set for RuleID
                buildGUONSSByRule(RuleID);
                if ( m_OpList.size() > 0)
                {
                    // 2. Create Enhancement String Descriptions
50         buildOpDescriptions();
                }
            }
        }
    }

```

```

// 3. Convert All Opnums in Enhancement Input Params
to GUONS
    buildTranslatedOpKeyIn();
    // 4. Reduce Duplicate GUONS (with same Description)
5    reduceSameDescriptionOps();
    // 5a. Generate OpKeysOut Information
    buildOpKeysOut();
    // 5b. Reduce GUONS that are NOT INPUT to any other
Non End GUON.
10    //    Dependent on 5a.
    reduceErrorGuons();
    // 6a. Perform Topology Sorting
    doTopologySort();
    // 6b. Perform "Node Coupling" Optimizations
15    doNodeCoupling();
    // 7. Allocate UNIR Data Structure Key Index
    allocateUNIROutputData();
    // 8. For every input parameter, match the appropriate
Index from step 7
20    propagateIndexToInput();
    }
    // 9. Create InstructionSet
    buildInstructionSet();
    // 10. Save InstructionSet in Hash Table with RuleID as Key
25    m_Instructions.put(new Integer(RuleID), m_TmplInstructions);
    // 11. Build FieldIndex Vector & send to Merger (if Target is
DB, not pipe)
    if (!m_IsTargetPipe)
    {
30        buildFieldIndex();
    }
    Log.Instance().writeLog("FlowManager.computeRuleFlows",
LogEvent.LOG_DEBUG,
        "Done Compute Flow for Rule "+RuleID);
35    }
    catch (Exception E)
    {
        Debug.writeStackTrace("FlowManager.computeRuleFlows", E);
    }
40    }

//*****
// Build an initial set of EnhanceOpNodes representing all Enhancements
per some
45    // rule. Set some additional info on every EnhanceOpNode, such as
GathererID and
    // m_IsLast Flag.
    // !!!! IT IS ASSUMED THAT Enhancements (From PERSISTENCE) is loaded
sorted by
50    // Field/Opnum !!!

```

```

//*****
protected void buildGUONSByRule( int RuleID )
{
    m_IsTargetPipe = false ;
5    m_TargetNode    = 1 ;
    m_TargetGatherer = 0 ;
    m_TargetIS       = 0 ;

    m_OpList = new EnhanceOpSList ();
10    Nodes NodeList = CEM.instance().getNodes();
        ISSs  ISSList = CEM.instance().getISSs();

    // Loop on All Enhancements. For Every Enhancement that matches the
current
15    // Rule, create an EnhanceOpNode object and add it to the OpArray. After
    // loop is done, Sort the OpArray.
    // OpArray will sort on ?
    Enhancements EnhancementList = CEM.instance().getEnhancements();
    Array OpArray = new Array();
20    if (EnhancementList != null)
    {
        for(Enumeration E = EnhancementList.elements();
E.hasMoreElements(); )
        {
25            Enhancement EnhancementItem = (
(Enhancement) (E.nextElement()) );
            if (EnhancementItem.getRuleID() == RuleID)
            {
                //printDiagnosticsPre(Enhancement);
30                EnhanceOpNode EnhancementNode = new EnhanceOpNode(
EnhancementItem );
                //printDiagnosticsPost(EnhancementNode);
                OpArray.add( EnhancementNode );
            }
35        }

        Sorting.sort( OpArray );

    // build m_OpList from OpArray
40    if (OpArray.size() > 0)
    {
        String LastField = "";
        Rules RuleList = CEM.instance().getRules();
        Rule CurRule = RuleList.getRuleByRuleID(RuleID);
45        // If Rule Not Enabled, m_OpList will be empty
        if ( (CurRule != null) && (CurRule.getEnabled()) )
        {
            // Get the AISM ID of this Flow
            int ISID = CurRule.getTriggerID();

```



```

// Set Target Pipe Boolean, Target Node and other Pipe
variables
    m_TargetIS = CurRule.getTargetID();
    if (m_TargetIS > 0)
    {
        m_IsTargetPipe = true ;
        IS TargetIS = ISsList.getISByISID(
m_TargetIS );
        if (TargetIS != null)
        {
            m_TargetGatherer =
TargetIS.getGathererID();
            Node N =
NodesList.getNodeByTypeID(Util.Comm.Node.NODE_TYPE_GATHERER, m_TargetGatherer);
            m_TargetNode = N.getNodeID();
        }
    }

// Get AISM Information ()
    IS I = ISsList.getISByISID( ISID );
    if (I != null)
    {
        int ctr = 0 ;
        // Build First entry into m_OpList from thr
        int GathererID = I.getGathererID();
        int ISMID = I.getISMID();
        Node N =
NodesList.getNodeByTypeID(Util.Comm.Node.NODE_TYPE_GATHERER, GathererID);
        int NodeID = N.getNodeID();

        // Add GUON 0 with data from the Rule Header
        EnhanceOpNode OpZero = new EnhanceOpNode
(RuleID, "", (short)0, ISID, 0, null, null);
        OpZero.setGathererID(GathererID);
        OpZero.setNodeID(NodeID);
        OpZero.setISMID(ISMID);
        // setting GUON to the original index will guarantee
        unique value.
        OpZero.setGUON(ctr);
        OpZero.setLastFlag(false) ;
        m_OpList.add ( OpZero );
        ctr++;

        // Loop on OpArray, build the rest of m_OpList
        EnhanceOpNode PriorEnhanceOp = null ;
        for(Enumeration E = OpArray.elements();
E.hasMoreElements(); )
    {

```

```

EnhanceOpNode CurrentEnhancement = (
    EnhanceOpNode) (E.nextElement());
    // Check if it belongs to The Rule
    if ( ( CurrentEnhancement != null ) &&
5   (CurrentEnhancement.getRuleID() == RuleID ) )
    {
        ctr++;
        // set aux data, add to list
        I = ISsList.getISByISID(
10   CurrentEnhancement.getISID());
        if (I!=null)
        {
            GathererID = I.getGathererID();
            ISMID = I.getISMID();
15   N =
            NodesList.getNodeByTypeID(Util.Comm.Node.NODE_TYPE_GATHERER, GathererID);
            NodeID = N.getNodeID();

20   CurrentEnhancement.setGathererID(GathererID);

            CurrentEnhancement.setNodeID(NodeID);

            CurrentEnhancement.setISMID(ISMID);
25   // set GUON to the original index will guarantee
            unique value.

            CurrentEnhancement.setGUON(ctr);

30   CurrentEnhancement.setLastFlag(false);

            // Check if Field changed. If
            Field has changed, this means that
            // the previous EnhanceOpNode
35   was the last for that field.
            String CurrentField =
            CurrentEnhancement.getFieldID();
            if ( (
40   !CurrentField.equals>LastField) ) && (PriorEnhanceOp != null) )
            {
                // Field Has Changed, prior
                Op was last for a Field

                PriorEnhanceOp.setLastFlag(true);
45   // Set Target Field of
                previous EnhanceOpNode

                PriorEnhanceOp.addOutFields>LastField);
50
            }

```

```

// Save for Last Step Per Field
Logic
PriorEnhanceOp =
CurrentEnhancement ;
5      LastField = CurrentField ;

// Add To List
m_OpList.add (
CurrentEnhancement );
10

Log.instance().writeLog("FlowManager.buildGUONSByRule",
LogEvent.LOG_DEBUG, CurrentEnhancement.toString() );

15      }
      else
      {
          Log.instance().writeLog(new
LogError("buildGUONSByRule: No IS found for "+CurrentEnhancement.getISID() ));
20      }
    }
    // End Of Loop Logic. Same as if Field has
changed.
25      if (PriorEnhanceOp != null)
      {
          PriorEnhanceOp.setLastFlag(true) ;
          PriorEnhanceOp.addOutFields(LastField);
      }
30      }
    }
}

35  //*****
// Print Diagnostics on the passed Enhancement.
// Currently Not Used.
//*****
protected void      printDiagnosticsPre(Enhancement EnhancementItem)
40  {
    try
    {
        byte[] buf = EnhancementItem.getOpKeyIn();
        String InputList = "[" ;
        Arguments a = new Arguments ( buf );
        45      for (Enumeration EK = a.elements(); EK.hasMoreElements(); )
        {
            short value =
((UNIRShort)EK.nextElement()).getShortValue();
50      InputList += Short.toString(value)+" ";

```

```

        value      =
        ((UNIRShort)EK.nextElement()).getShortValue();
        InputList += Short.toString(value)+" ";
    }
5      InputList += "{";
        Log.instance().writeLog("FlowManager.buildGUONSByRule",
                                LogEvent.LOG_DEBUG, "PersistenceItem:
        "+InputList);

10      }
        catch(Exception e)
        {
        }
    }

15      //*****
      // Print Diagnostics on the passed EnhanceOpNode.
      // Currently Not Used.
      //*****

20      protected void      printDiagnosticsPost(EnhanceOpNode
EnhancementNode)
      {
          try
          {
25              byte[] buf = EnhancementNode.getOpKeyIn();
              String InputList = "{" ;
              Arguments a = new Arguments ( buf );
              for (Enumeration EK = a.elements(); EK.hasMoreElements(); )
              {
30                  short value =
                    ((UNIRShort)EK.nextElement()).getShortValue();
                    InputList += Short.toString(value)+" ";
                    value      =
                    ((UNIRShort)EK.nextElement()).getShortValue();
35                    InputList += Short.toString(value)+" ";
              }
              InputList += "}";
              Log.instance().writeLog("FlowManager.buildGUONSByRule",
                                      LogEvent.LOG_DEBUG, "Struct 1:
40          "+InputList);

              InputList = "{" ;
              for (Enumeration EK =
EnhancementNode.getParsedOpKeyIn().elements(); EK.hasMoreElements(); )
45              {

                  int value = ( (Integer)EK.nextElement() ).intValue();
                  InputList += Integer.toString(value)+" ";
                  value      = ( (Integer)EK.nextElement() ).intValue();
50                  InputList += Integer.toString(value)+" ";

```

```

    }
    InputList += " ";
    Log.instance().writeLog("FlowManager.buildGUONSByRule",
        LogEvent.LOG_DEBUG, "Struct 2:
5    "+InputList);

    }
    catch(Exception e)
10    {
    }

    }

    /*******
15    // ABSTRACT
    // Convert the m_OpList to an Instruction Set the Gatherer Can
    Understand
    // (a set of UNIRInstructions)
    //
    // Description
20    // Translate m_OpList to sets of Enhancement Instructions. Divide the
    Enhancements
    // by Different Gatherer sequences, wrap each segment with a Receive
    Instruction
25    // and a Send Instruction.
    /*******

    /*******
    // Helper method to add an instruction of type Send
    /*******
30    void AddSendStep(int StepCtr, int NodeID, int ISID, int
    ISMID, int GathererID)
    {
    UNIRInstruction u = new UNIRInstruction (m_RuleID,
35    StepCtr,

    UNIRInstruction.INSTRUCTION_SEND,

    NodeID,
40    0, ISID, ISMID,
    null, null,

    GathererID );

    Log.instance().writeLog("FlowManager.buildInstructionSet",
    LogEvent.LOG_DEBUG,
45    "Adding Instruction "+u);
    m_TmpInstructions.pushBack(u);

    }

50    /*******

```

```

// Helper method to add an receive of type Receive
//*****
void AddReceiveStep(int StepCtr, int NodeID, int ISID, int
ISMID, int GathererID)
5      {
        UNIRInstruction u = new UNIRInstruction (m_RuleID,

        StepCtr,

10      UNIRInstruction.INSTRUCTION_RECEIVE,

        NodeID,

        0, ISID, ISMID,

        null, null,

        GathererID);
15      Log.getInstance().writeLog("FlowManager.buildInstructionSet",
        LogEvent.LOG_DEBUG,

        "Adding Instruction "+u);
        m_TapInstructions.pushBack(u);

20      }

//*****
// Helper method to add a receive of type Enhance
//*****
void AddEnhanceStep(int StepCtr, int NodeID, int ISID, int
ISMID, int GathererID,
        int FunctionID, byte[] inParam, byte[]
outParam)
30      {
        UNIRInstruction u = new UNIRInstruction (m_RuleID,

        StepCtr,

        UNIRInstruction.INSTRUCTION_ENHANCE,

35      NodeID,

        FunctionID, ISID,

        ISMID,

        inParam, outParam, GathererID);
40      Log.getInstance().writeLog("FlowManager.buildInstructionSet",
        LogEvent.LOG_DEBUG,

        "Adding Instruction "+u);
        m_TapInstructions.pushBack(u);

45      }

//*****
// Helper method to add an receive of type Pipe Give
//*****

```

```

        void AddGiveStep(int StepCtr, int NodeID, int ISID, int
ISMID, int GathererID,
                                int FunctionID, byte[] inParam, byte[]
outParam)
5          {
            UNIRInstruction u = new UNIRInstruction (m_RuleID,
                StepCtr,
10         UNIRInstruction.INSTRUCTION_GIVE,
                                NodeID,
                                FunctionID, ISID,
                ISMID,
                inParam, outParam, GathererID);
            Log.instance().writeLog("FlowManager.buildInstructionSet",
LogEvent.LOG_DEBUG,
                "Adding Instruction "+u);
            m_TmpInstructions.pushBack(u);
20          }

        //*****

25      protected void      buildInstructionSet()
        {
            Log.instance().writeLog("FlowManager.buildInstructionSet",
LogEvent.LOG_DEBUG,
                "Build Instruction Set for Rule
30      "+m_RuleID);
            EnhanceOpNode PriorOp = null ;
            int PriorNodeID = -1;
            m_TmpInstructions = new SList ();
            int StepCtr = 0;
35      for (SListIterator I = m_OpList.begin(); !I.atEnd(); I.advance())
        {
            EnhanceOpNode Op1 = (EnhanceOpNode)I.get();
            if ( !Op1.getLastFlag() )
            {
40                if (Op1.getNodeID() != PriorNodeID)
                    {
                        // Add Send Instruction to close prior node
sequence
                        if (PriorNodeID != -1)
45                            {
                                AddSendStep(StepCtr,
                                    Op1.getNodeID(), Op1.getISID(),
                                    Op1.getISMID(),
                                    PriorOp.getGathererID());
50                                StepCtr++;

```

```

        }
        // Add Recieve Instruction to start current node
sequence
        AddReceiveStep(StepCtr, Opl.getNodeID(),
5  Opl.getISID(),
                                Opl.getISMID(),
        Opl.getGathererID());

        StepCtr++;
        PriorNodeID = Opl.getNodeID();
10      }
        PriorOp = Opl ;
        // Add Enhancement Step
        AddEnhanceStep(StepCtr, Opl.getNodeID(),
        Opl.getISID(), Opl.getISMID(),
15      Opl.getGathererID(),
        Opl.getFunctionID(),

        Opl.getOpKeyInAsArray(), Opl.getOpKeyOutAsArray());
        StepCtr++;
20      }
    }
    // The Last Step, Add Send Instruction to CEM
    // Check global Pipe flag. If this Flow Target is some Pipe ISM,
    // Set the instructions Type to INSTRUCTION_GIVE and set correct Node
25  if (PriorOp != null)
    {
        if (m_IsTargetPipe)
        {
            // if target is in a different node, add a send and recieve
30      if (m_TargetNode != PriorOp.getNodeID())
            {
                AddSendStep(StepCtr, m_TargetNode, 0, 0,
                PriorOp.getGathererID());

                StepCtr++;
35      AddReceiveStep(StepCtr, 0, 0, 0,
                m_TargetGatherer);

                StepCtr++;
            }
            // adding the final Give instruction. Dont forget Binding
40  Information
                byte[] bindingInfo =
                buildFISMBindingInformation();
                AddGiveStep(StepCtr, m_TargetNode, m_TargetIS, 0,
                                m_TargetGatherer, 0,
45  bindingInfo, null);
            }
        else
        {
            // Send Data to CEM

```



```

        AddSendStep(StepCtr, Util.Comm.Node.NODE_TYPE_CEM,
        C, 0, PriorOp.getGathererID());
    }
    StepCtr++;
5      }

    }

    /*******
    // doTopologySort()
10    // m_OpList can be seen as a Directed Acyclic Graph (DAG), and an
    algorithm
    // (Introduction To Algorithms, Chapter 23) to sort the graph is applied.
    The
    // result of this procedure is a list where the order of operations is
15    correct,
    // meaning that all enhancements whose output is used as input happen
    before the
    // enhancements that use them.
    /*******

20    protected void visitDFS(EnhanceOpNode Op)
    {
        Op.setColor(EnhanceOpNode.GRAY);
        Op.setDFSDiscoverTime(m_TimeCtr);
25        m_TimeCtr++;
        // For Each Adjacent Vertex (All Output GUONS)
        for (Enumeration AdjVert = Op.getOutputGUONS();
        AdjVert.hasMoreElements(); )
        {
30            int OutGuon =
            ((Integer)AdjVert.nextElement()).intValue();
            EnhanceOpNode OutOp = m_OpList.findGuon(OutGuon);
            // Explore Edges Each (Guon, OutGuon) pair is an
            edge
35            if ( ( OutOp != null) && (OutOp.getColor() ==
            EnhanceOpNode.WHITE) )
            {
                visitDFS(OutOp);
                OutOp.setPredecessor(Op);
40            }
        }
        Op.setColor(EnhanceOpNode.BLACK);
        Op.setDFSFinishTime(m_TimeCtr);
        m_TimeCtr++;
45        m_OpTmpList.pushFront(Op);
    }

    //-----
50

```

```

protected void doTopologySort()
{
    // Create Temporary List to Store Ordered Ops.
    m_OpTmpList = new EnhanceOpSList ();
5
    m_TimeCtr = 0 ;
    // For each vertex
    for (int i=0;i<m_OpList.size();i++)
    {
10        EnhanceOpNode Op1 = (EnhanceOpNode)m_OpList.at(i);
        // if not used in DFS yet (Color White)
        if ( Op1.getColor() == EnhanceOpNode.WHITE)
        {
            visitDFS( Op1 );
15        }
    }

    // Temporary List is now the correct one, move it to normal list
    m_OpList = new EnhanceOpSList ( m_OpTmpList );
20
    //*****
    // doNodeCoupling()
25    // Node Coupling attempts to move enhancements that share the same
    gatherer closer
    // in sequence, as a way to reduce traffic between gatherers. This is
    done without
    // damaging the basic order of the enhancements established earlier in
30 the
    // Topology Sort.
    //*****
    //-----
35    // Find if inserting Op at Pos will cause sort order
    conflict
    protected boolean noHarmToSort(EnhanceOpSList OpList, int
    pos, EnhanceOpNode Op)
40    {
        for (int i=pos;i<OpList.size();i++)
        {
            EnhanceOpNode OpToCompare =
            (EnhanceOpNode)OpList.at(i);
45            // Check if Op has input that uses Output
            of OpToCompare
            // Loop on all inputs of Op, see if the Guon value
            matches OpToCompare.Guon
            for (Enumeration EK = Op.getInIterator();
50 EK.hasMoreElements(); )

```

```

        {
            OpInputParam InParam =
                (OpInputParam)EK.nextElement();

            if ( InParam.getGUON() ==
5      OpToCompare.getGUON() )
            {
                // No insertion Possible
                return false ;
            }
10      }
        }
        // No Conflict Found, Insertion Allowed
        return true ;
    }

15      //-----
        // find last position of first same node sequence in
        OpList
        protected int findSameNodePos(EnhanceOpSList OpList,
20      EnhanceOpNode Op)
        {
            int NodeID = Op.getNodeID();
            boolean StartSeqFound = false ;
            int ctr = 0;
            int ResultVal = -1;
25      for (SListIterator I = OpList.begin(); (!I.atEnd());
                I.advance())
            {
                EnhanceOpNode OpToCompare =
30      (EnhanceOpNode)I.get();

                if (NodeID == OpToCompare.getNodeID())
                {
                    StartSeqFound = true ;
                }
                else
35      {
                    if ( StartSeqFound )
                    {
                        ResultVal = ctr ;
                        break; // get out of loop, we found
40      what we're looking for
                    }
                }
                ctr++;
45      }

            return ResultVal ;
        }

50      //-----

```

```

protected void doNodeCoupling()
{
    // Create Temporary List to Store Ordered Ops.
5   m_OpTmpList = new EnhanceOpSList ();

    for (SListIterator I = m_OpList.begin(); (!I.atEnd()); I.advance())
    {
        EnhanceOpNode Op1 = (EnhanceOpNode)I.get();
10        // Add Op1 to the Temporary List, either at the end (No Node
        Optimization Possible)
        // or Insert in the middle (Node Optimization, Couple Nodes)
        int n1 = findSameNodePos(m_OpTmpList, Op1);
        if (n1 == -1)
15        {
            // No Node Optimizations possible, add to end
            m_OpTmpList.pushBack(Op1);
        }
        else
20        {
            if ( noHarmToSort(m_OpTmpList, n1, Op1) )
            {
                // Insertion OK, insert to group same node Ops
                together
25                m_OpTmpList.insert(n1, Op1);
            }
            else
            {
                // Insertion would hurt the input->output order,
30                cannot inset, add to end
                m_OpTmpList.pushBack(Op1);
            }
        }
    }
35    // Temporary List is now the correct one, move it to normal list
    m_OpList = new EnhanceOpSList ( m_OpTmpList );
}

40    //*****
    // Create an image of the Data to be produced by this Rule. This
    allocates output
    // positions for every Enhancement. These positions are used later to
45    update the
    // input params.
    //*****
    protected void allocateGNIROutputData()
50    {

```

```

        UNIRAllocation = new SList();
        // For Each OP
        for (int i=0;i<m_OpList.size();i++)
        {
5           EnhanceOpNode Opl = (EnhanceOpNode)m_OpList.at(i);
            if (Opl != null)
            {
                // For Each Output Parameter
                for (Enumeration EK = Opl.getOutputKeys();
10            EK.hasMoreElements(); )
                {
                    Integer Key = (Integer)EK.nextElement(); // First
                    // is key, 2nd is guon
                    GuonKeyParam NewKey = new GuonKeyParam(Key, new
15            Integer(Opl.getGUON() ) );
                    if (NewKey!=null)
                    {
                        UNIRAllocation.pushBack(NewKey);
                    }
20                }
            }
        }

25        //*****
        // propagateIndexToInput()
        // Based on the Unir Allocation, match every input parameter with the
        index in the
30        // UNIRAllocation.
        //*****
        *****

        protected int findIndexInAllocation ( OpInputParam KeyIn)
65        {
            GuonKeyParam CurInKey = KeyIn.getGuonKey();
            for (int i=0;i<UNIRAllocation.size();i++)
            {
                GuonKeyParam KeyInAllocation =
40            (GuonKeyParam)UNIRAllocation.at(i);
                if (CurInKey.equals(KeyInAllocation))
                {
                    return i ;
                }
45            }
            return -1 ;
        }

        //---
50

```

```

protected void propagateIndexToInput()
{
    // For Each OP
    for (SListIterator I = m_OpList.begin(); (!I.atEnd()); I.advance())
    {
        EnhanceOpNode Opl = (EnhanceOpNode)I.get();
        // For Each Input Parameter
        for (Enumeration EK = Opl.getInputIterator(); EK.hasMoreElements();
    {
        OpInputParam InParam = (OpInputParam)EK.nextElement();
        // Get index of pair in UNIRAllocation
        int n = findIndexInAllocation(InParam);
        if (n == -1)
        {
            // TBC, Error
        }
        InParam.setIndex(n + FLOW_OFFSET);
    }
    Log.instance().writeLog("FlowManager.propagateIndexToInput",
    LogEvent.LOG_DEBUG,
        "Propagated Input Op: " +Opl);
}

//*****
// For every Enhancement
// Build List of Enhancements that need the output of some enhancement
as Input.
//
// This information is needed later to weed out Non productive
enhancements and
// To perform the Topological sort on m_OpList
//*****
protected void buildOpKeysOut()
{
    // Op1 Loop : Loop on All Enhancements in m_OpList
    for (SListIterator I = m_OpList.begin(); (!I.atEnd()); I.advance())
    {
        EnhanceOpNode Opl = (EnhanceOpNode)I.get();
        // Op2 Loop : Loop on All Enhancements in m_OpList
        // Compare all Enhancements to Opl (Except itself)
        for (SListIterator J = m_OpList.begin(); (!J.atEnd());
        J.advance())
        {
            EnhanceOpNode Op2 = (EnhanceOpNode)J.get();
            // No need to compare Op to itself
            if (Opl!=Op2)

```

```

(
    // If Op1 participates as input in Op2, add Op2 to
    list of
    // Enhancements that need Op1 output. This list is
5    managed in Op1.
    for (Enumeration EK = Op2.getIterator();
    EK.hasMoreElements(); )
    {
        OpInputParam InParam =
10    (OpInputParam)EK.nextElement();
        if (InParam.getGUON() == Op1.getGUON())
        {
            Op1.addOutputParam(InParam.getKey(), Op2.getGUON());
15
        }
    }
}
}
}
20
}

//*****
*****
// Reduce Duplicates. No need for duplication of process if another
25 enhancement
// already creates the needed Output. Keep one of the duplicates, delete
the other.
// Pass thru all input keys of all enhancements, replacing any reference
to the
30 // Deleted GUON with the Surviving GUON.
//*****
*****
protected void reduceSameDescriptionOps()
{
35 for (int i=0; i<m_OpList.size(); i++)
{
    EnhanceOpNode Op1 = (EnhanceOpNode)m_OpList.at(i);
    // Move from end to Current Pos because we might delete an object
    for (int j=m_OpList.size()-1; j>i; j--)
40 {
        EnhanceOpNode Op2 = (EnhanceOpNode)m_OpList.at(j);
        if ( Op1.getDescription().equals( Op2.getDescription() ) )
        {
            Log.instance().writeLog("FlowManager.reduceSameDescriptionOps",
45 LogEvent.LOG_DEBUG,
            "Reducing Description :
            "+Op2.getFieldID()+" ":"+
            Op2.getGUON()+" ":"+
50

```

```

Op2.getDescription() );
// Merge outfields of deleted guon with surviving
guon
5      Op1.addOutFields(Op2.getOutFields() );
      int newGUON = Op1.getGUON();
      int oldGUON = Op2.getGUON();
      m_OpList.remove(j);
      // GUON Replacement Pass
10     for (int k=0;k<m_OpList.size();k++)
        {
            EnhanceOpNode Op3 =
(EnhanceOpNode)m_OpList.at(k);
            Op3.swapInputKeyGUON(oldGUON, newGUON);
15         }
        // Mark for Garbage Collection
        Op2 = null ;
    }
}

//*****
*****
25     // Reduce Error Guons, that do not produce an output needed as input by
any other
    // guon. They can be identified if they are not "Last" and the output
list is empty.
    //*****
30     *****
    protected void reduceErrorGuons()
    {
        if (m_OpList.size() > 0)
        {
35         for (int i=m_OpList.size()-1;i>=0;i--)
            {
                EnhanceOpNode Op1 = (EnhanceOpNode)m_OpList.at(i);
                // If Not Last Guon per some field...
                if (!Op1.getLastFlag())
40                 {
                    // Check if any Outputs are defined
                    Enumeration ChkOutput =
Op1.getOutputGUONKeys();
                    if (!ChkOutput.hasMoreElements())
45                     {
                        Log.getInstance().writeLog("FlowManager.reduceErrorGuons",
LogEvent.LOG_DEBUG,
                        "Reducing Bad Enhancement
;"+Op1.getGUON()+" "+Op1.getDescription());

```



```

// No Output Elements for this
Guon, delete it
m_OpList.remove(i);
// Mark for Garbage Collection
5      OpI = null ;
      }
    }
  }

10      }

//*****
// Build Translated OpKeyIn for every Enhancement
// !!!! IT IS ASSUMED THAT Enhancements (From PERSISTENCE) are loaded
sorted by
// Field/Opnum !!!
//*****
20      protected void buildTranslatedOpKeyIn()
      {
        for (SListIterator I = m_OpList.begin(); (!I.atEnd()); I.advance())
        {
          EnhanceOpNode Op = (EnhanceOpNode)I.get();
25          if (Op != null)
          {
            Op.translateOpKeyIn(m_OpList);

            Log.instance().writeLog("FlowManager.buildTranslatedOpKeyIn",
30      LogEvent.LOG_DEBUG, Op.toString() );
          }
        }
      }

35      //*****
// Build a Description for every Enhancement
// !!!! IT IS ASSUMED THAT Enhancements (From PERSISTENCE) is loaded
sorted by
// Field/Opnum !!!
//*****
40      protected void buildOpDescriptions()
      {
        for (SListIterator I = m_OpList.begin(); (!I.atEnd()); I.advance())
        {
          EnhanceOpNode Op = (EnhanceOpNode)I.get();
          if (Op != null)
          {
50          string S = getDescription( Op );

```

```

        Log.getInstance().writeLog("FlowManager.buildOpDescriptions",
        LogEvent.LOG_DEBUG,
                                "Op Description :

5      "+Op.getFieldID()+":"+Op.getGUON()+": " +S);
        }
    }

10      //*****
      *****
      // get Textual representation of the first EnhanceOpNode Input Key
      //*****
      *****
15      protected String getKeyText(EnhanceOpNode Op)
      {
          String ResultVal = "" ;
          try
          {
20              if (Op != null)
              {
                  Arguments args = new Arguments( Op.getOpKeyIn() );
                  int n =
((UNIRShort)args.get(ArgKeys.KEY_KEY_ID)).getShortValue();
25              ResultVal = Integer.toString(n);
              }
          }
          catch (Exception e) {}
          return ResultVal ;
30      }

      //*****
      *****
      // Build a Textual description of a regular EnhanceOpNode Op
35      //
      // the methods "parseArgs" and "getDescription" are used in a recursive
      manner
      //*****
      *****
40      protected String parseArgs(EnhanceOpNode Op)
      {
          String s = "" ;
          try
          {
45              // Loop On All Op Key In
              Arguments a = new Arguments ( Op.getOpKeyIn() );
              for (Enumeration EK = a.elements(); EK.hasMoreElements(); )
              {
                  // Add Description of Input Param OpNum

```

```

short OpValue = ((UNIRShort)
EK.nextElement()).getShortValue();
EnhanceOpNode CurrentOp = m_OpList.findOp(OpValue,
Op.getFieldID());
5      if ( !Op.equals(CurrentOp) )
      {
          s += getDescription(CurrentOp)+",";
          // Add Description of Input Param Key
          short KeyValue = ((UNIRShort)
10      EK.nextElement()).getShortValue();
          s += Integer.toString(KeyValue);
          // If More Param, delimit by comma
          if ( EK.hasMoreElements() )
          {
15              s+=", " ;
          }
          else
          {
20              Log.instance().writeLog("FlowManager.parseArgs", LogEvent.LOG_ERROR,
              "Corruption Detected in Data (Illegal Recursion)");
              throw ( new Exception() );
          }
25      }
      }
      catch(Exception EA)
      {
          s = "" ;
30      }
      return s ;
      }

//*****
35      *****
      // Build a Textual description of the EnhanceOpNode Op
      //
      // the methods "parseArgs" and "getDescription" are used in a recursive
      manner
40      //*****
      *****
      protected String getDescription(EnhanceOpNode Op)
      {
          String s = Op.getDescription();
45          // If No description was built yet, build it and save for next
          time
          if (s==null)
          {
              // if last Step for some field, just parse
50          if (Op.getLastFlag())

```

```

        {
            s = parseArgs(Op);
        }
        else
5           if (Op.getOpNum() == 0)
            {
                // First Op, has no function
                s = Op.getISID() + "." + getKeyText(Op);
            }
10          else
            {
                // Function, Must show Func ID + Input
Parameters.
                s = Op.getISID() + "." + Op.getFunctionID()
15          + "{";
                s += parseArgs(Op);
                s += " ";
            }
            //---- Save for next time
20          Op.setDescription(s);
            //----
        }
        return s ;
    }
25
    /*******
    *
    // Scan All Rules, Loop On All Instructions, send matching instructions
    to Gatherer
30    /*******
    *****
    public void setInstructionsByGatherer( int GathererID, GathererController GC
    )
    {
35        Rules RuleList = CEM.instance().getRules();
        if (RuleList != null)
        {
            // For Each Rule
            for(Enumeration E1 = RuleList.elements(); E1.hasMoreElements(); )
40            {
                // Build Flow for specific Rule
                Rule CurrentRule = ( Rule) (E1.nextElement());
                if (( CurrentRule != null ) && (CurrentRule.getEnabled()
== true) )
45            {
                int RuleID = CurrentRule.getRuleID();
                setInstructionsByRuleByGatherer( GathererID,
RuleID, GC );
            }
50        }
    }

```

```

    }
    }

    //*****

5   *
    // Send All Instructions per Rule/ per Gatherer to the Gatherer
    // Thru the GathererController
    //*****

    *
10   synchronized public void setInstructionsByRuleByGatherer( int GathererID,
    int RuleID, GathererController GC )
    {
        SList TmpInstructions = (SList)m_Instructions.get(new
        Integer(RuleID));
15       Vector v = new Vector(); // Hold UNIRInstruction belonging to
        some Gatherer
        // Loop on All instructions of a Rule, Add Gatherer Matches to v,
        Send v to GC
        for (Enumeration EOps = TmpInstructions.elements();
20       EOps.hasMoreElements();)
        {
            UNIRInstruction Op = (UNIRInstruction)EOps.nextElement();
            if (Op.getGathererID() == GathererID)
            {
25                 v.addElement(Op);
            }
        }
        try
        {
30             Log.instance().writeLog("FlowManager.setInstructionsByRuleByGatherer",
            LogEvent.LOG_DEBUG, "Disable Flow "+RuleID+ " on " + GathererID);
            GC.disableFlow(RuleID);
            if (v.size() > 0)
35             {

                Log.instance().writeLog("FlowManager.setInstructionsByRuleByGatherer",
                LogEvent.LOG_DEBUG, "Upload Flow "+RuleID+ " on " + GathererID);
                GC.sendFlowData(RuleID, new UNIRInstruction(v) );
40             }

            Log.instance().writeLog("FlowManager.setInstructionsByRuleByGatherer",
            LogEvent.LOG_DEBUG, "Enable Flow "+RuleID+ " on " + GathererID);
            GC.enableFlow(RuleID);
        }
45     }
    catch(Exception e)
    {
    }
    }
50

```

```

//*****
*
// Pack the Binding Information between current Flow to the input
// parameters of the target Pipe ISM.
5 // This information is added to the Give instruction and sent to the
// Gatherer where it is used to pass data from the output UNIR of
// the current flow to the input parameters of the target flow.
// RETURNS
// a byte array containing Binding information in the following format:
10 // pairs of shorts made of the index of the data in the output UNIR,
followed
// by the key value of the target input parameter.
//*****
*
15 byte[] buildPISMBindingInformation()
{
    try
    {
        Arguments args = new Arguments();
        Vector v = new Vector();
        // Loop on All Enhancements
        for (SListIterator I = m_OpList.begin(); (!I.atEnd());
20 I.advance())
        {
            EnhanceOpNode Op = (EnhanceOpNode)I.get();
            if (Op != null)
            {
                int Index = 0; // Position of data in
                Output Unir
30 short Key = 0; // Key ID of data in
                Output Unir
                int inKey = 0; // Key ID of data in
                Input Parameters
                // the binding process will
                // "take the UNIR in the Index position from the
35 flow data and put
                // in the PISM input parameter in inKey"

                // One to many relation: One input to Many output
                // DEFINE INPUT : Get Index & key of first input
40 key
                Enumeration EK = Op.getIterator();
                if ( EK.hasMoreElements() )
                {
                    OpInputParam InParam =
45 (OpInputParam)EK.nextElement();

                    Index = InParam.m_KeyIndex ;
                    Key = (short)InParam.m_KeyID;

50
                }
            }
        }
    }
}

```

```

// DEFINE OUTPUT(S) Loop on all Out Values
for (Enumeration EFlds =
Op.getOutFields().elements(); EFlds.hasMoreElements();)
{
5      String FieldID =
      (String)EFlds.nextElement();

      Log.instance().writeLog("FlowManager.buildPISMBindingInformation",
LogEvent.LOG_INFO,
10      "Got Pipe
      Binding: "+ FieldID +" ; "+ Index +" ; "+ Key);
      // in PISM, FieldID is a String
      Representation of the KeyID
      // defined as string for
15      compatability with Fields (see buildFieldIndex)
      inKey = Integer.valueOf( FieldID
      ).intValue();

      // Add to Arguments;
20      args.put(ArgKeys.ARGKEY_ENHANCEMENT_OP_NUM, new UNIRShort((short)Index));
      args.put(ArgKeys.KEY_KEY_ID, new
      UNIRShort((short)inKey));

25      }
    }
    return args.toByteArray();
  }
  catch(Exception e)
  {
30      Log.instance().writeLog("FlowManager.buildPISMBindingInformation",
LogEvent.LOG_ERROR,
35      "Error in
      Binding Info, "+e);

      Debug.writeStackTrace("FlowManager.buildPISMBindingInformation", e);

40      }
      return null ;
    }

45      //*****
      *
      // Builds a vector of all DB fields that participate in a specific flow.
      // (The current flow, in m_OpList)
      // This list is built of FieldIndexStruct elements and sent to the
50      Merger.

```

```

        // The Merger uses this information in processing incoming UNIRs from
this
        // flow.
        // The vector is a cross reference between the DB FieldID and the
5 position
        // the data to be put in the field from the incoming UNIR for this flow.
        //*****
*
        void buildFieldIndex()
10
        {
            Vector v = new Vector();
            // Loop on All Enhancements
            for (SListIterator I = m_OpList.begin(); (!I.atEnd()); I.advance())
            {
15                EnhanceOpNode Op = (EnhanceOpNode)I.get();
                if (Op != null)
                {
                    int Index = 0;
                    short Key = 0;
20                    Enumeration EK = Op.getInIterator();
                    // Get Index & key of first input key
                    if (EK.hasMoreElements() )
                    {
                        OpInputParam InParam =
25 (OpInputParam)EK.nextElement();

                        Index = InParam.m_KeyIndex ;
                        Key = (short)InParam.m_KeyID;

                    }
30                    // Loop on all Fields Out, add all fields in list to
vector with
                    // Index & key stored earlier.
                    for (Enumeration EFlds =
Op.getOutFields().elements(); EFlds.hasMoreElements(); )
35 {
                        String FieldID =
                        (String)EFlds.nextElement();

                        v.addElement(new FieldIndexStruct(FieldID,
Index, Key));
40
                    }
                }
            }
            if (v.size() > 0)
            {
45                // Call Merger
                Merger m = CEM.instance().getMerger();
                Log.instance().writeLog(new LogDebug("Vector:\n"+v));
                m.addRule(m_RuleID, v);
            }
50            else

```



```

        // Call Merger
        Merger m = CEM.instance().getMerger();
        Log.instance().writeLog(new LogDebug("Deleting Rule"+m_RuleID));
5         m.deleteRule(m_RuleID);
    }
}

//*****

10 // Get All Enhancements with matching FieldID
// Return Vector of affected Rules
//*****

15 public Vector getRulesWithField(String FieldID)
{
    Log.instance().writeLog("FlowManager.getRulesWithRuleField",
LogEvent.LOG_DEBUG, "Get Enhancements with Field "+FieldID);
    Vector ResultVal = new Vector();
20    OrderedSet AffectedRules = new OrderedSet(false);
    Vector ToDel = new Vector();
    Enhancements EnhancementList = CEM.instance().getEnhancements();
    for(Enumeration E = EnhancementList.elements();
E.hasMoreElements(); )
25    {
        Enhancement CurrentEnhancement = (
Enhancement)(E.nextElement());
        if ( ( FieldID.equals( CurrentEnhancement.getFieldID() ) ) )
        {
30            AffectedRules.add( new Integer (
CurrentEnhancement.getRuleID() ) );
        }
    }

35    // Set Result Val
    for(Enumeration AR = AffectedRules.elements();
AR.hasMoreElements(); )
    {
40        Integer RuleIDObj = (Integer)AR.nextElement();
        ResultVal.addElement(RuleIDObj); // add Rule to Result Set
    }
    return ResultVal ;
}

45 //*****

// Delete All Enhancements with matching FieldID, RuleID from List
//*****

```

```

public void deleteRulesWithField(Vector RuleList, String FieldID)

    for (int i=0;i<RuleList.size();i++)
5      {
        int RuleID =
        ((Integer) (RuleList.elementAt(i))).intValue();
        deleteRulesWithRuleField(RuleID, FieldID);
      }
10

//*****
*
15 // Delete All Enhancements with matching RuleID, FieldID
// Return Vector of affected Rules
//*****
*

public Vector deleteRulesWithRuleField(int RuleID, String FieldID)
20
{
    Log.instance().writeLog("FlowManager.deleteRulesWithRuleField",
    LogEvent.LOG_DEBUG, "Delete Enhancements with Field "+FieldID);
    Vector ResultVal = new Vector();
    OrderedSet AffectedRules = new OrderedSet(false);
25    Vector ToDel = new Vector();
    Enhancements EnhancementList = CEM.instance().getEnhancements();
    for(Enumeration E = EnhancementList.elements();
    E.hasMoreElements(); )
30    {
        Enhancement CurrentEnhancement = (
        (Enhancement) (E.nextElement()));
        if ( ! FieldID.equals( CurrentEnhancement.getFieldID() ) ) ;
        { (RuleID == CurrentEnhancement.getRuleID() ) ; }
35        /* (RuleID == CurrentEnhancement.getRuleID() ) || (RuleID
        == -1) ) )
        {
            ToDel.addElement(CurrentEnhancement);
        }
40    }
    for (int i=0;i<ToDel.size();i++)
    {
        Enhancement CurrentEnhancement = (Enhancement)ToDel.elementAt(i);
        AffectedRules.add( new Integer ( CurrentEnhancement.getRuleID() )
45    );

    Log.instance().writeLog("FlowManager.deleteRulesWithRuleField",
    LogEvent.LOG_DEBUG, "Deleting "+CurrentEnhancement);
    try { CurrentEnhancement.delete(); } catch (Exception e){}
50
}

```

```

        // Set Result Val
        for(Enumeration AR = AffectedRules.elements();
AR.hasMoreElements(); )
5      {
            Integer RuleIDObj = (Integer)AR.nextElement();
            ResultVal.addElement(RuleIDObj); // add Rule to Result Set
        }
        return ResultVal ;
10    }

    //*****
    *
15    // deleteRuleEnhancements
    //*****
    *

    public void deleteRuleEnhancements(int RuleID)
    {
20        Vector ToDel = new Vector();
        Enhancements EnhancementList = CEM.instance().getEnhancements();
        for(Enumeration E = EnhancementList.elements();
E.hasMoreElements(); )
        {
25            Enhancement CurrentEnhancement = (
(Enhancement)(E.nextElement()) );
            if ( CurrentEnhancement.getRuleID() == RuleID )
            {
                ToDel.addElement(CurrentEnhancement);
30            }
        }
        for (int i=0;i<ToDel.size();i++)
        {
            Enhancement CurrentEnhancement = (Enhancement)ToDel.elementAt(i);
35            try
            {
                CurrentEnhancement.delete();
            }
            catch (Exception e)
40            {
                Log.instance().writeLog(new LogDebug("Delete Rule
Enhancement Failed... "+e));

                Debug.writeStackTrace("FlowManager.deleteRuleEnhancements", e);
45            }
        }
    }

    //*****
50

```

```

// Delete All Enhancements with matching ISID
// If All enhancements per rule are deleted, disable the Rule
// If Trigger is Deleted, Delete Rule
// Return Vector of affected Rules
5 //*****

public Vector deleteRulesWithIS(int ISID)
{
    Log.instance().writeLog(new LogDebug("Delete Enhancements for "+ISID));
10     Vector ResultVal = new Vector();
    OrderedSet AffectedRules = new OrderedSet(false);
    OrderedSet AffectedRuleFields = new OrderedSet(false);
    Vector ToDel = new Vector();
    // Loop on all Rules
15     Rules RuleList = CEM.instance().getRules();
    for(Enumeration E1 = RuleList.elements(); E1.hasMoreElements(); )
    {
        // Compare Trigger ID to the ISID being deleted
        Rule CurrentRule = (Rule)(E1.nextElement());
20         int RuleID = CurrentRule.getRuleID();
        if (CurrentRule.getTriggerID() == ISID)
        {
            // yes. found a match. Add to Affected List
            // and delete the whole Flow
25             AffectedRules.add( new Integer ( RuleID ) );
            ToDel.addElement(CurrentRule);
            deleteRuleEnhancements(RuleID);
        }
    }
30     // Delete from Collection
    for (int i=0;i<ToDel.size();i++)
    {
        Log.instance().writeLog(new LogDebug("Delete Rule with Trigger
35         "+ISID));
        Rule CurrentRule = (Rule)ToDel.elementAt(i);
        try { CurrentRule.delete(); } catch (Exception e){}
    }

    // Loop on All Enhancements, find Rule/Field pairs affected by IS
40     removal
    Enhancements EnhancementList = CEM.instance().getEnhancements();
    for(Enumeration E2 = EnhancementList.elements();
    E2.hasMoreElements(); )
    {
45         Enhancement CurrentEnhancement = (
        (Enhancement)(E2.nextElement()));
        if (CurrentEnhancement.getISID() == ISID)
        {
50             Log.instance().writeLog(new LogDebug("Adding Rule Field to
            Affected List "+

```

```

CurrentEnhancement.getRuleID()+" : "+
CurrentEnhancement.getFieldID() );
5      AffectedRuleFields.add(new Pair(new Integer
      (CurrentEnhancement.getRuleID()),

CurrentEnhancement.getFieldID() ) );
      }
10      }

      for(Enumeration E3 = AffectedRuleFields.elements();
      E3.hasMoreElements(); )
      {
15          Pair CurrentPair = (Pair)(E3.nextElement());
          // Delete All Enhancements for this Rule & FieldID

          deleteRulesWithRuleField(((Integer)CurrentPair.first).intValue(),
          (String)CurrentPair.second);
20          AffectedRules.add( (Integer)CurrentPair.first );
          }

          // Set Result Val
          for(Enumeration AR = AffectedRules.elements();
25          AR.hasMoreElements(); )
          {
              Integer RuleIDObj = (Integer)AR.nextElement();
              ResultVal.addElement(RuleIDObj); // Add Rule to Result Set
          }
30      }

      // Return Affected Rules
      return ResultVal ;
      ;

35      //*****
      // Delete from the given Rule all the Fields that have Enhancements using
      // at least 1 from the given set of Trigger Input Keys .
      //*****
40      ;

      public OrderedSet deleteRuleFieldsWithTriggerKeys (int RuleId, Vector
      TriggerKeys )
      {
          Log.instance().writeLog(new LogDebug
45          ("deleteRuleFieldsWithTriggerKeys "+"RuleId=" + RuleId+
              ", TriggerKeys="+ TriggerKeys));

          int OpNum = 0;
          OrderedSet AffectedRule = new OrderedSet(false);
          OrderedSet AffectedFields = new OrderedSet(false);
50

```

```

        if ( TriggerKeys==null || TriggerKeys.size()==0)
            return null;
        Log.Instance().writeLog(new LogDebug ("Add Rule to Affected List "+
RuleId));
5         AffectedRule.add(new Integer(RuleId));

        //Get all Enhancements
        Enhancements EnhancementList = CEM.Instance().getEnhancements();
        for(Enumeration E = EnhancementList.elements(); E.hasMoreElements(); )
10         {
            Enhancement CurrentEnhancement = ( Enhancement)(E.nextElement());
            if (CurrentEnhancement.getRuleID()==RuleId)
                // Search InputKeys for Trigger ones (OpNum = 0)
                {
15                 EnhanceStruct CurrentEnhanceStruct = new
EnhanceStruct(CurrentEnhancement);
                    Vector ParsedOpKeyIn = CurrentEnhanceStruct.getParsedOpKeyIn();
                    ///
                    int CurrentOpNum;
20                 int CurrentOpKey;
                    for (int i=0; i<ParsedOpKeyIn.size(); i+=2)
                    {
                        CurrentOpNum =
((Integer)ParsedOpKeyIn.elementAt(i)).intValue();
25                         CurrentOpKey =
((Integer)ParsedOpKeyIn.elementAt(i+1)).intValue();
                        if ( CurrentOpNum==OpNum && TriggerKeys.contains (new
Integer(CurrentOpKey)) )
                        {
30                             Log.Instance().writeLog(new LogDebug("Adding Rule Field to
Affected List "+
CurrentEnhancement.getRuleID()+" : "+
CurrentEnhancement.getFieldID() ));
35                             AffectedFields.add( CurrentEnhanceStruct.getFieldID());
                                break; //pass to next Enhancement
                            }
                        } // for over ParsedOpKeyIn
40                    } //if (CurrentEnhancement.getRuleID()==RuleId)

                //loop over Enhancements
                for(Enumeration E = AffectedFields.elements(); E.hasMoreElements(); )
                {
45                    String FieldId = (String)E.nextElement();
                        // Delete All Enhancements for this Rule & FieldID
                        deleteRulesWithRuleField(RuleId, FieldId);
                }
                return AffectedRule;
50            } //deleteRuleFieldsWithTriggerKeys

```

```

//*****
// Delete in all Rules all Fields including an Enhancement from the given
// ISID with FunctionId equal to one from the given set.
//*****
5 public OrderedSet deleteRulesFieldsWithISFunctions (int ISID , Vector
ChangedFunctions)
{
    Log.instance().writeLog(new LogDebug
10 ("deleteRulesFieldsWithISFunctions "+ "ISID="+ISID+
        ", ChangedFunctions="+ChangedFunctions));
    Vector ToDel= new Vector();
    OrderedSet AffectedRules = new OrderedSet(false);
    OrderedSet AffectedFields = new OrderedSet(false);
15
    if ( ChangedFunctions==null || ChangedFunctions.size()==0)
        return null;
    Enhancements EnhancementList = CEM.instance().getEnhancements();
    for(Enumeration E = EnhancementList.elements(); E.hasMoreElements(); )
20 {
        Enhancement CurrentEnhancement = ( Enhancement)(E.nextElement());
        if (CurrentEnhancement.getISID() == ISID &&
            ChangedFunctions.contains( new Integer(CurrentEnhancement.getFunctionID()) ) )
        {
25 //ToDel.addElement(CurrentEnhancement);
            Log.instance().writeLog(new LogDebug("Adding Rule Field to
Affected List "+
                CurrentEnhancement.getRuleID()+" : "+
30 CurrentEnhancement.getFieldID() ));
            AffectedFields.add(new Pair(new Integer
                (CurrentEnhancement.getRuleID()),
                CurrentEnhancement.getFieldID() ));
35
        }
    }
    for(Enumeration E2 = AffectedFields.elements(); E2.hasMoreElements(); )
    {
40 Pair CurrentPair = ( Pair)(E2.nextElement());
        // Delete All Enhancements for this Rule & FieldID

        deleteRulesWithRuleField(((Integer)CurrentPair.first).intValue(),
            (String)CurrentPair.second);
45 Log.instance().writeLog(new LogDebug("Add Rule to Affected
List "+CurrentPair.first));
        AffectedRules.add( (Integer)CurrentPair.first );
    }
50 return AffectedRules;

```

```

//deletesRulesFieldsWithISFunctions

//*****

5
// In all Rules find all Enhancements which met the following conditions:
//   ISID = <given ISID>, FunctionID = <one from the given list>.
// For all these Enhancements the Functionid is replaced with new one .
// The "new" Id is placed immediately after the matching "old" Id in the
// 2-nd input parameter.
10
//*****

    public OrderedSet editEnhancementsFunctionIDWithFunctionID(int ISID ,
    Vector AffectedFunctionIds)
        throws Exception
15
    {
        Log.instance().writeLog(new LogDebug
        ("editEnhancementsFunctionIDWithFunctionID " + "ISID="+ISID+
        ", AffectedFunctionIds="+AffectedFunctionIds ));
        OrderedSet AffectedRuleSet = new OrderedSet(false);
20
        if ( AffectedFunctionIds==null || AffectedFunctionIds.size()==0)
        {
            Log.instance().writeLog(new LogDebug
            ("editEnhancementsFunctionIDWithFunctionID: ", "There are no changed function
25 IDs" ));
            return null;
        }
        Enhancements EnhancementList = CEM.instance().getEnhancements();
        for(Enumeration E = EnhancementList.elements(); E.hasMoreElements(); )
30
        {
            Enhancement CurrentEnhancement = ( Enhancement)(E.nextElement());
            for (int i=0;i<AffectedFunctionIds.size();i+=2)
            {
                if (CurrentEnhancement.getFunctionID()==
35 ((Integer)AffectedFunctionIds.elementAt(i)).intValue())
                {
                    int OldFuncId = CurrentEnhancement.getFunctionID();
                    int NewFuncId =
                    ((Integer)AffectedFunctionIds.elementAt(i)).intValue();
40
                    Log.instance().writeLog(new LogDebug ("Replace FunctionId
                    "+OldFuncId +
                    "with " + NewFuncId+ " in Enhancement
                    "+CurrentEnhancement));
                    CurrentEnhancement.edit();
50
                    CurrentEnhancement.setFunctionID(NewFuncId);
                    CurrentEnhancement.update();
                    Log.instance().writeLog(new LogDebug("Add Rule to Affected
                    List "+CurrentEnhancement.getRuleID()));
                    AffectedRuleSet.add (new Integer
                    (CurrentEnhancement.getRuleID()));

```



```

        }
    }

5     return AffectedRuleSet;
    }
} //class

//*****
10 // Internal Helper Classes to support Flow computations
//*****

//*****
,
15 //
//*****
*
class EnhanceOpSList extends SList
{
20 //*****
,
//
25 //*****
*
public EnhanceOpSList()
{
    super();
30 }

//*****
,
35 //
//*****
*
public EnhanceOpSList(int n)
40 {
    super(n);
}

45 //*****
,
//
50 //*****
,

```

```

public EnhanceOpSlist(EnhanceOpSlist s)
{
    super(s);
}

//*****
//
//*****

public EnhanceOpNode findOp (int OpValue, String OpField)
{
    for (int i=0;i<size();i++)
    {
        EnhanceOpNode Op = (EnhanceOpNode)at(i);
        if (Op.getOpNum() == OpValue)
        {
            // Field does not matter if we are searching for
            // OpNum 0
            if ( !OpField.equals( Op.getFieldID() ) ) {}
            {OpValue == 0}
            {
                return Op ;
            }
        }
        return null ;
    }
}

//*****
//
//*****

public EnhanceOpNode findGuon (int GuonValue)
{
    for (int i=0;i<size();i++)
    {
        EnhanceOpNode Op = (EnhanceOpNode)at(i);
        if (Op.getGHON() == GuonValue)
        {
            return Op ;
        }
        return null ;
    }
}

```

```

    //*****
    *
    5
    //*****
    *
    // Object that represents a single Input Parameter of a single
    10 Enhancement
    // Operation.
    // Every Enhancement should maintain a set of OpInputParam
    //*****
    *
    15 class OpInputParam
    {
        // Original OpNum of Enhancement owning the Input Param
        public int m_OpNum = 0 ;
        // ArgKey Value of the Output Parameter
    20 public int m_KeyID = 0 ;
        // GUON value of Enhancement owning the Input Param
        public int m_OpGUON = 0 ;
        // Index of Input Parameter, Defined from the "UNIR Allocation"
        public int m_KeyIndex = 0 ;
    25
        //*****
        // Constructors
        //*****
        public OpInputParam( int OpNum, int KeyID, int GUON)
    30 {
            m_OpNum = OpNum ;
            m_KeyID = KeyID ;
            m_OpGUON = GUON ;
            m_KeyIndex = 0 ;
    35 }

        //*****
        // Get Set
        //*****
    40 public int getGUON() {return m_OpGUON;}
        public void setGUON(int Value) {m_OpGUON = Value;}
        public int getKey() {return
m_KeyID;}
        public void setKey(int Value) {m_KeyID = Value;}
    45
        public GuonKeyParam getGuonKey() {return new GuonKeyParam(new
Integer(m_KeyID), new Integer(m_OpGUON));}

        public int getIndex() {return m_KeyIndex;}
    50 public void setIndex(int Value) {m_KeyIndex = Value;}

```

```

    }
    //*****

5
    //*****

    //
10
    //*****

    class GuonKeyParam extends Pair
    {
    public int m_GUON = 0 ;
15
    public int m_Key = 0 ;

        public GuonKeyParam ()
        {
            super();
20
        }

        public GuonKeyParam (Integer Key, Integer GUON)
        {
            super (Key, GUON);
25
        }

        public int hashCode()
        {
            Integer n1 = (Integer)first;
            Integer n2 = (Integer)second;
30
            return ( (n1.intValue() << 16)+ n2.intValue() );
        }

        public boolean equals(Pair pair)
35
        {
            Integer n1 = (Integer)first;
            Integer n2 = (Integer)pair.first;
            Integer n3 = (Integer)second;
            Integer n4 = (Integer)pair.second;
40
            return ( ( n1.intValue() == n2.intValue() ) && (
n3.intValue() == n4.intValue() ) );
        }

45
    }

    //*****

    // Object that represents a single Output Parameter of a single
50
    Enhancement Operation.

```

```

// Every Enhancement should maintain a set of OpOutputParam
//*****

*
class OpOutputParam
5   {
    // Set of Key/GUON pairs, elements are JGL Pairs
    public OrderedSet m_KeyGUONList = new OrderedSet ( false ); // No
    Duplicates
    // Set of Keys, elements are Integers
10   public OrderedSet m_KeyList      = new OrderedSet ( false ); // No
    Duplicates
    // Set of GUONs, elements are Integers
    public OrderedSet m_GUONList      = new OrderedSet ( false ); // No
    Duplicates
15
    //*****
    *
    // Set Set
20
    //*****
    *
    public Enumeration getGUONs()          (return m_GUONList.elements());
    public Enumeration getKeys()           (return
25   m_KeyList.elements());
    public Enumeration getGUONKeys()       (return m_KeyGUONList.elements());

    //*****
    *
30   // FFD:
    //public Keys getGUONKeys(GUON);
    //public GUONs getKeyGUONs(Key);

35   //*****
    //

    //*****
    public void addKey(int Key, int GUON)
40   {
        Pair KeyGuon = new GuonKeyParam (new Integer(Key), new
        Integer (GUON) );
        m_KeyGUONList.add(KeyGuon);
        m_KeyList.add(new Integer(Key));
45   m_GUONList.add(new Integer(GUON));
    }

    //*****
50   //

```

```

//*****
    public byte[] getOpKeyOutAsArray()
    {
5         try
        {
            Arguments args = new Arguments();
            for (Enumeration EK = getKeys(); EK.hasMoreElements(); )
            {
10                 Integer outKey = (Integer)EK.nextElement();
                    args.put(ArgKeys.KEY_KEY_ID, new
UNIRShort((short)outKey.intValue()) );

            }
15         return args.toByteArray();
        }
        catch (Exception e){}

        return null ;
20    }

//*****

25    )

//*****
    *
    // Object that represents a single Enhancement Operation.
    // Flow Manager should maintain a set of EnhanceOpNode
30    //*****
    *

    class EnhanceOpNode extends EnhanceStruct
    {
        // Unique Identifier of the Op, Created at Runtime. "Global
35    Unique Op Num"
        int    m_GUON = 0 ;
        // String Description of the Operation, Created "Recursively" at Runtime.
        String m_Description = null ;

40        // OrderedSet version of EnhanceStruct.OpKeysIn. Each element is
        of type OpInputParam
        OrderedSet    m_TranslatedOpKeysIn    = new OrderedSet ( false ); // No
        Duplicates
        Vector        m_TranslatedOpKeysInVector = new Vector();

45        // List of Fields that are the "End Product" of this GUON. null
        in most objects, has value only for
        // "End" Objects. Because of Optimizations, one GUON can feed
        more than one field. Elements are Strings
50        // that were the FieldID of Original Enhancement.

```

```

OrderedSet m_OutFields = new OrderedSet( false );

// Object that manages the lists of output param keys of this Enhancement
OpOutputParam m_OpKeysOut = new OpOutputParam();

5
// ISM, Gatherer & Node ID of this (derived from EnhanceStruct.ISID)
int m_ISMID      = 0;
int m_GathererID = 0;
int m_NodeID     = 0;

10
// Flag if this GUON was originally the "last" (Last OpNum for a FIELDID)
boolean m_IsLast = false ;

// The list of EnhanceOpNodes can be viewed as a Directed Acyclic
15 Graph (DAG), with the GUONS as vertices
// and the pairs of (GUON n, GUON m using output of n as input)
as edges.
// The following are Variables Used in a TOPOLOGICAL SORT using
Depth First Search (DFS).

20
final static int WHITE = 0;
final static int GRAY = 1;
final static int BLACK = 2;
public int m_ColorStatus = WHITE;
public EnhanceOpNode m_Pi = null; // Predecessor Node. If
25 null, then root (Several roots possible)
public int m_DiscoverTime = 0 ; // Start Time of DFS on this
GUON
public int m_FinishTime = 0; // End Time of DFS
on this GUON (and adjacency list)

30
//////////
// Constructors
//////////
public EnhanceOpNode( Enhancement E)
35
{
    super (E);
}

public EnhanceOpNode(int RuleID, String FieldID, short OpNum, int ISID,
40 int FunctionID, byte[] OpKeyIn,
Vector ParsedOpKeyIn
{
    super (RuleID, FieldID, OpNum, ISID, FunctionID, OpKeyIn,
45 ParsedOpKeyIn );
}

//////////
// Operations
50
//////////

```

```

//*****
*
5      // Create a list of OpInputParam in m_TranslatedOpKeysIn from the
pair
      // arguments in m_OpKeyIn. Add the corresponding GUON for every
OpNum
      // in the OpKeyIn.
10
//*****
*
      public void translateOpKeyIn(EnhanceOpSList OpList)
      {
15          try
          {
              Arguments a = new Arguments ( getOpKeyIn() );
              for (Enumeration EK = a.elements(); EK.hasMoreElements();
              {
20                  short OpValue = ((UNISShort)
EK.nextElement()).getShortValue();
                  short KeyValue = ((UNISShort)
EK.nextElement()).getShortValue();
25                  OpInputParam tmp = null ;
                  EnhanceOpNode CurrentOp = OpList.findOp(OpValue,
m_FieldID);
                  tmp = new OpInputParam (OpValue, KeyValue,
CurrentOp.getGUON() );
30                  m_TranslatedOpKeysIn.add(tmp);
                  m_TranslatedOpKeysInVector.addElement(tmp);
              }
              catch(Exception EA){}
          }
35
//*****
*
      // pass thru m_TranslatedOpKeysIn, replacing occurrences of
40      oldGUON with newGUON
//*****
*
      public void SwapInputKeyGUON (int oldGUON, int newGUON)
      {
45          for (Enumeration EK = m_TranslatedOpKeysIn.elements();
EK.hasMoreElements(); )
          {
              OpInputParam InParam = (OpInputParam)EK.nextElement();
50              if (InParam.getGUON() == oldGUON)

```



```

        {
            InParam.setGUON(newGUON);
        }
    }

5
    }

//*****
*
10    // Add Op to Lists of output of current EnhanceOpNode

//*****
*

    public void addOutputParam    (int Key, int GUON)
15    {
        m_OpKeysOut.addKey(Key, GUON);
    }

20
//*****
*

// Get Set

//*****
25
*

    public int getGUON()                {return m_GUON;}
    public void setGUON(int Value)      {m_GUON = Value;}
    public int getGathererID()          {return
m_GathererID;}
30    public void setGathererID(int Value) {m_GathererID = Value;}
    public int getNodeID()              {return
m_NodeID;}
    public void setNodeID(int Value)    {m_NodeID = Value;}
    public boolean getIsLast()          {return m_IsLast ;}
35    public void setIsLast(boolean value) {m_IsLast = value ;}
    public String getDescription()      {return m_Description ;}
    public void setDescription(String value) {m_Description = value ;}
    public int getISMID()               {return
m_ISMID;}
40    public void setISMID(int Value)    {m_ISMID =
Value;}

//*****
45
*

// Topology Sort

//*****
*

```

```

        public int getColor() {return
m_ColorStatus;}
        public void setColor(int Value) {m_ColorStatus =
Value;}
5      public void setDFSDiscoverTime(int Value) {m_DiscoverTime =
Value;}
        public void setDFSFinishTime(int Value) {m_FinishTime =
Value;}
10      public void setPredecessor(EnhanceOpNode Value) {m_Pi=Value;}

        //*****
        *
        public void addOutFields(String FieldID)
15      {m_OutFields.add(FieldID);}
        public void addOutFields(OrderedSet FieldList) {m_OutFields
= m_OutFields.union(FieldList);}
        public OrderedSet getOutFields() {return
m_OutFields;}
20      public Enumeration getIterator() {return
m_TranslatedOpKeysInVector.elements();}
        public Enumeration getOutputGUONs() {return
m_OpKeysOut.getGUONs();}
        public Enumeration getOutputGUONKeys() {return
25      m_OpKeysOut.getGUONKeys();}
        public Enumeration getOutputKeys() {return
m_OpKeysOut.getKeys();}

30      //*****
        *
        // Convert OpKeyIn structure to a byte array for the
UNIRInstruction
35      //*****
        *
        public byte[] getOpKeyInAsArray()
        {
40          try
          {
              Arguments args = new Arguments();
              for (Enumeration EK = getIterator();
EK.hasMoreElements(); )
45              {
                  OpInputParam InParam =
(OpInputParam)EK.nextElement();
                  args.put(ArgKeys.ARGKEY_ENHANCEMENT_OP_NUM,
new UNIRShort((short)InParam.m_KeyIndex));
50

```

```

                                args.put(ArgKeys.KEY__KEY__ID, new
UNIRShort((short)InParam.m_KeyID));

                                }
5                                return args.toByteArray();
                                }
                                catch(Exception e){}

                                return null ;
10                                }

                                /*******
                                *
                                // Convert OpKeyOut structure to a byte array for the
15                                UNIRInstruction

                                /*******

                                *
20                                public byte[] getOpKeyOutAsArray()
                                {
                                    return m_OpKeysOut.getOpKeyOutAsArray();
                                }

25                                /*******
                                *
                                // String representation of this Enhancement

30                                /*******
                                *

                                public String toString()
                                {
                                    String s = m_GUON + " [" ;
35
                                    for (Enumeration EK = getInIterator(); EK.hasMoreElements(); )
                                    {
                                        OpInputParam InParam = (OpInputParam)EK.nextElement();
                                        s += "(" + InParam.m_OpGUON + " ; " +
40                                InParam.m_OpNum + " ; " + InParam.m_KeyIndex + " ; " + InParam.m_KeyID+" )";

                                    }

                                    s += " ] . [" ;
45

                                    for (Enumeration EK2 = m_OutFields.elements();
EK2.hasMoreElements(); )
                                    {
                                        String OutFld = (String)EK2.nextElement();

```

```

        s += "("+OutField+"":

    }
    s += "], {" :
5
    for (Enumeration EK3 = getOutputGUONKeys();
EK3.hasMoreElements(); )
    {
        Pair GuonKey = (Pair)EK3.nextElement();
10
        s += "("+GuonKey+"":

    }

    s += "], " +
15
        m_ISMID + ":" + m_GathererID + ":" + m_NodeID ;
    return s ;
}

20
//*****

// to support special operation of slist

25
//*****

public boolean equals(Object obj)
{
    return ( m_GUON == ((EnhanceOpNode)obj).getGUON() );
30
}

//*****

35
// to support special operation of slist

//*****

public int hashCode()
40
{
    //byte v1 = (byte)(m_GUON & 0x000000ff) ;
    byte v1 = 0 ;
    int v2 = 0 ;
    if (!m_FieldID.equals(""))
45
        v2 = m_FieldID.hashCode() & 0x0000ffff ;
    return ( v1 << 24) + (v2 << 8) + m_OpNum ;
}

50
}

```

5
.....

CLAIMS

What is claimed is:

1. A system for tracking network session information, the system comprising:
 - 5 an information source module having a source information input and a standardized information output, a source information corresponds to network usage information, a standardized information corresponds to the network usage information transformed into a standard format;
 - a first program having at least a first standardized information input and
10 an enhanced data output, a first standardized information input corresponding to the standardized information, an enhanced data corresponding to the standardized data after at least a partial transformation, the at least partial transformation being defined according to a data record format;
 - 15 a second program having at least a first enhanced data input and a data record output, the first enhanced data corresponding to the enhanced data, a data record corresponding to the first enhanced data, the data record being formatted according to the data record format;
 - a database storing the data record; and
 - 20 wherein the second program merges duplicate data records that represent the same network usage information.

2. The system of claim 1 wherein the at least partial transformation
is defined from a data enhancement procedure, and wherein the data record
format includes a plurality of fields and wherein the data enhancement
5 procedure defines how the standardized information is to be transformed into
the plurality of fields of the data record format.

3. The system of claim 2 wherein the data enhancement procedure
includes at least a field enhancement wherein the field enhancement defines a
source for a predetermined field in the plurality of fields.

10 4. The system of claim 2 wherein the data enhancement procedure
includes at least a field enhancement wherein the field enhancement defines a
function to be applied to at least a portion of the standardized data.

5. The system of claim 2 wherein the data enhancement procedure
defines a plurality of field enhancements, wherein each field enhancement
15 defines network usage information to be stored in the plurality of fields.

6. The system of claim 5 wherein at least a first field enhancement
corresponds to an information source module that performs aggregation on at
least a portion of the network usage information to be stored in at least a first
field in the plurality of fields.

20 7. The system of claim 6 wherein the aggregation occurs by
aggregating packet flow information, the packet flow information corresponding

to a plurality of packets, where each packet in the aggregated packets has the same IP source address, destination address and port information.

8. The system of claim 5 wherein at least a first field enhancement corresponds to a set of information source modules that performs filtering and aggregation on at least a portion of the network usage information to be stored
5 in at least a first field in the plurality of fields.

9. The system of claim 5 wherein at least a first field enhancement corresponds to a set of information source modules that performs merging, filtering, and aggregation on at least a portion of the network usage information
10 to be stored in at least a first field in the plurality of fields.

10. The system of claim 5 wherein at least a first field enhancement corresponds to a set of information source modules that performs event notification and provisioning activation.

11. The system of claim 2 further comprising a second information
15 source module, the second information source module having a second source information input and a second standardized information output, a second source information corresponds to a second network information, a second standardized information corresponds to the second network information transformed into a standard format, and wherein the data enhancement
20 procedure includes a first definition of at least a first field in the plurality of fields being from the standardized information, and at least a second definition

of a second field in the plurality of fields being from the second standardized information.

12. The system of 11 further comprising a first information source and a quality of service information source, and wherein the information source
5 module receives the network usage information from the first information source, and wherein the second information source module receives the second network information from the quality of service information source, and wherein the first definition defines that a source IP address supplied by the first information source should be put into the first field, and wherein the second
10 definition defines a supplied by the DNS server should be put into the second field.

13. The system of claim 1 wherein the second program manages the first program and the information source module.

14. The system of claim 1 wherein the second program causes the
15 data record to be stored in the database.

15. The system of claim 1 wherein the information source module is configured to receive the network usage information from a predetermined network device.

16. The system of claim 1 wherein the at least partial transformation
20 includes policy-based data aggregation which defines how network usage data should be aggregated.

17. The system of claim 1 wherein the network usage information includes IP session data.

18. The system of claim 1 wherein the data format includes a plurality of fields including a source IP field, a destination IP field, a source
5 host field, a destination host field, a service type field, a date and time field, a duration field, a total number of bytes field, and a counter field.

19. The system of claim 1 further comprising a customer care and billing system coupled to the database, the customer care and billing system for accessing the database to generate a bill from the data record.

10 20. A network usage accounting system comprising:
an information source module coupled to receive network information from a network device;
a gatherer coupled to receive the network information source module, the gatherer for performing data enhancements on the network
15 information to create a plurality of data records;
a central database storing the plurality of data records; and
a central event manager coupled to receive the plurality of data records, the central event manager merging duplicate records in the plurality of data records, the duplicate records representing the same network
20 usage information.

21. The system of claim 20 wherein the information source module is configured to receive network information from a network device chosen

from the group of network devices consisting of a proxy server, a domain name service server, a firewall, a RADIUS server, and a router.

22. The system of claim 20 wherein the gatherer performs filtering and aggregation on the network information.

5 23. The system of claim 20 wherein the plurality of data records have a predefined data format comprising a plurality of fields, and wherein the data enhancements includes at least a first data field enhancement to enhance the network information to fill in the first data field.

10 24. The system of claim 23 wherein the first data field corresponds to a source IP address field and wherein the data enhancement includes extracting a source IP address value from the network information.

25. The system of claim 23 wherein the first data field corresponds to a domain name field and wherein the data enhancement includes requesting a domain name from a domain name service server.

15 26. A method of gathering and aggregating network usage information from a set of network devices, the system using at least a first program and a second program coupled in communications, the method comprising:

accessing network communications usage information;
20 filtering and aggregating the network communications usage information using the first program;

completing a plurality of data records from the filtered and aggregated
network communications usage information, the plurality of data
records corresponding to network usage by a plurality of users;
storing the plurality of data records; and
5 merging duplicate records in the plurality of data records.

27. The method of claim 26 wherein completing the plurality of
records includes accessing user account information.

28. The method of claim 26 wherein completing the plurality of
records includes for each data record determining a corresponding source IP
10 address, a corresponding domain name, a corresponding type of service used,
and a corresponding amount of time used.

29. The method of claim 26 wherein the system includes a third
program coupled in communications with at least the second program and
wherein completing the plurality of records includes accessing the third program
15 to determine network account information and including the network account
information in at least a first record in the plurality of records.

30. The method of claim 26 wherein merging the duplicate records
includes comparing a plurality of fields in the data records to identify data
records corresponding to the same network session and merging the
20 corresponding records.

31. The method of claim 26 wherein merging the duplicate records
includes automatically deleting a duplicate record.

32. The method of claim 26 further comprising using the second program to automatically update the filtering and aggregation performed by the first program.

33. A network usage tracking system comprising:

5 means for accessing network communications usage information;

means for filtering and aggregating the network communications usage information using the first program;

means for completing a plurality of data records from the filtered and aggregated network communications usage information, the plurality

10 of data records corresponding to network usage by a plurality of users;

means for storing the plurality of data records; and

means for merging duplicate records in the plurality of data records.

34. The network usage tracking system of claim 34 wherein the

15 means for completing the plurality of data records includes one or more networked computers running one or more programs.

35. The network usage tracking system of claim 34 wherein the means for storing the plurality of data records includes a relational database.

36. The network usage tracking system of claim 34 wherein the

20 means for storing the plurality of data records includes an object database.

37. A system for accounting for network usage comprising:

a plurality of gatherers coupled to receive network usage information,

each gatherer for performing data enhancements on the network

information to create a plurality of data records; and

a central event manager coupled to receive the plurality of data records,


5 the central event manager merging duplicate records in the plurality
of data records, the duplicate records representing the same network
usage information.

38. The system of claim 37 further comprising a customer care and

billing application for receiving the plurality of data records and generating

10 bills.

1/7



System 100
Fig. 1

Transaction Data From Multiple Network Nodes
Data Distillation

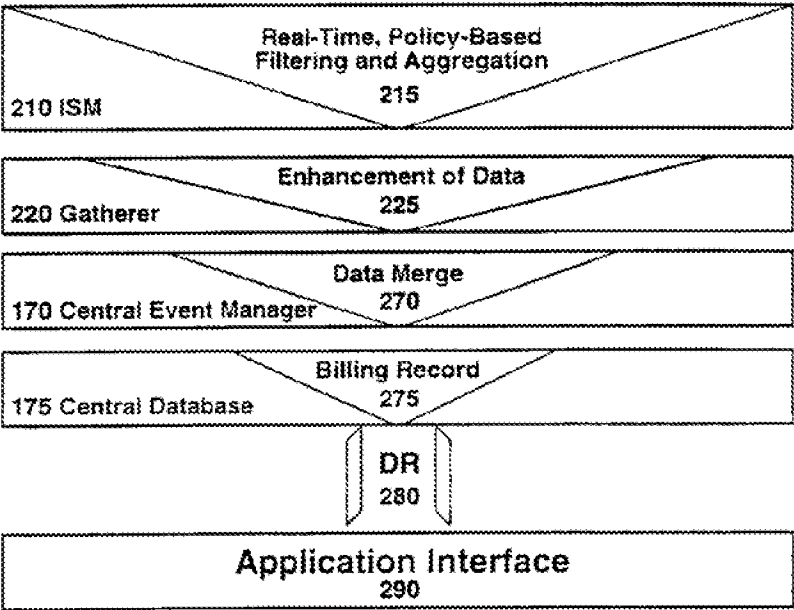


Fig. 2

3/7




Fig. 3

417

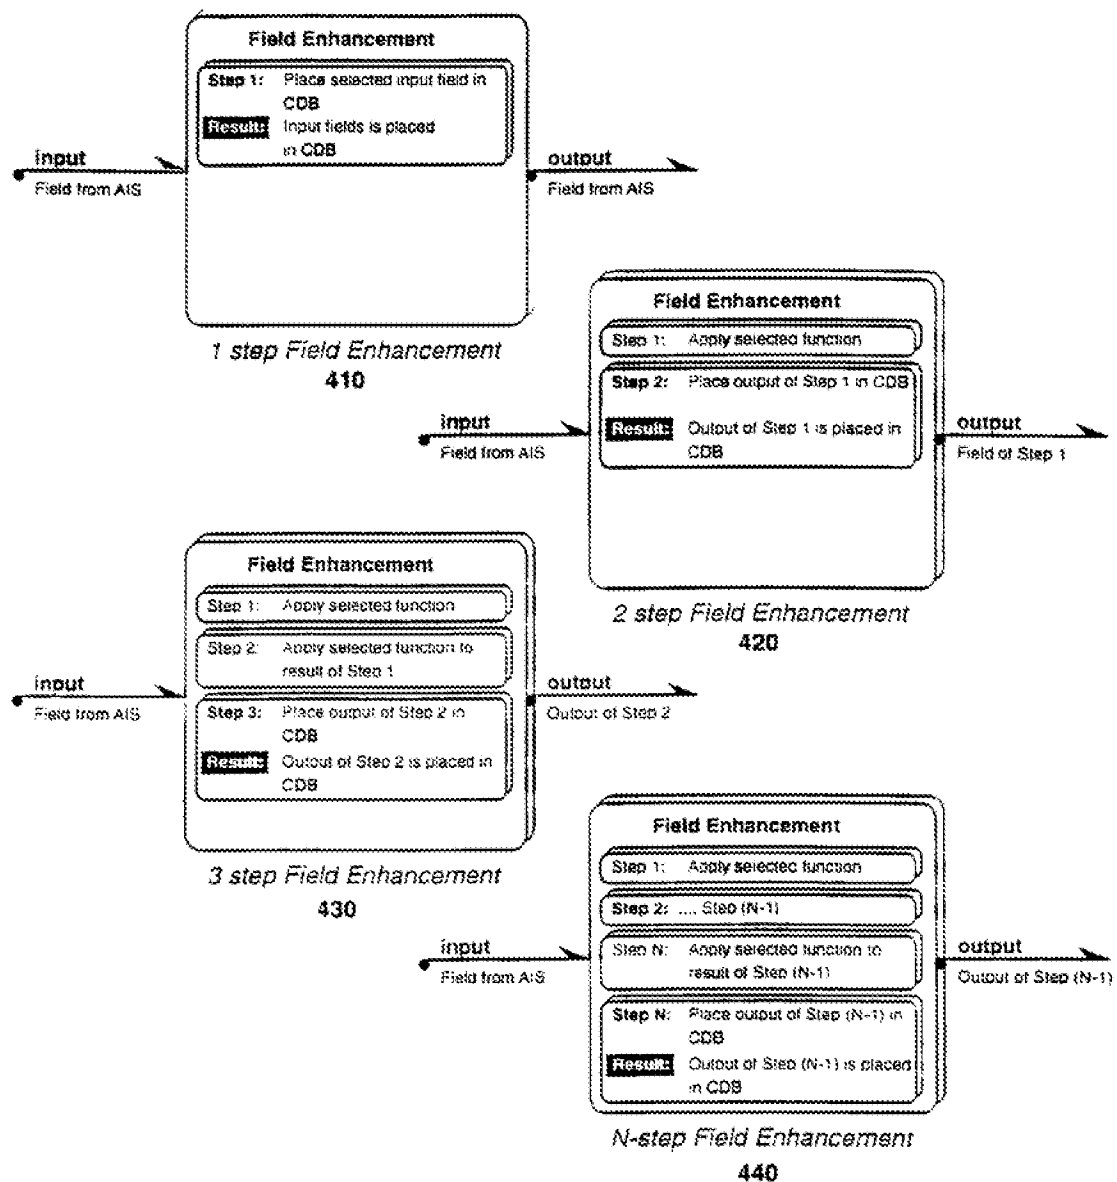


Fig. 4A

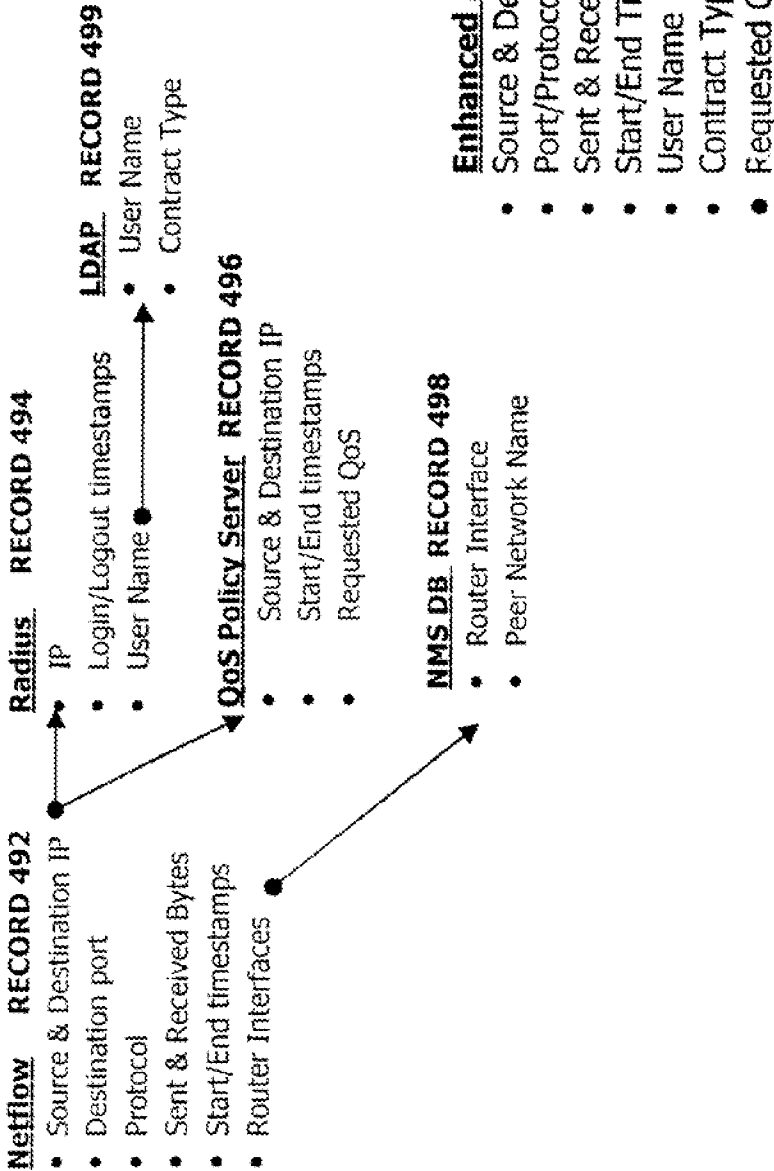


Fig. 4B

6/7




Fig. 5

777

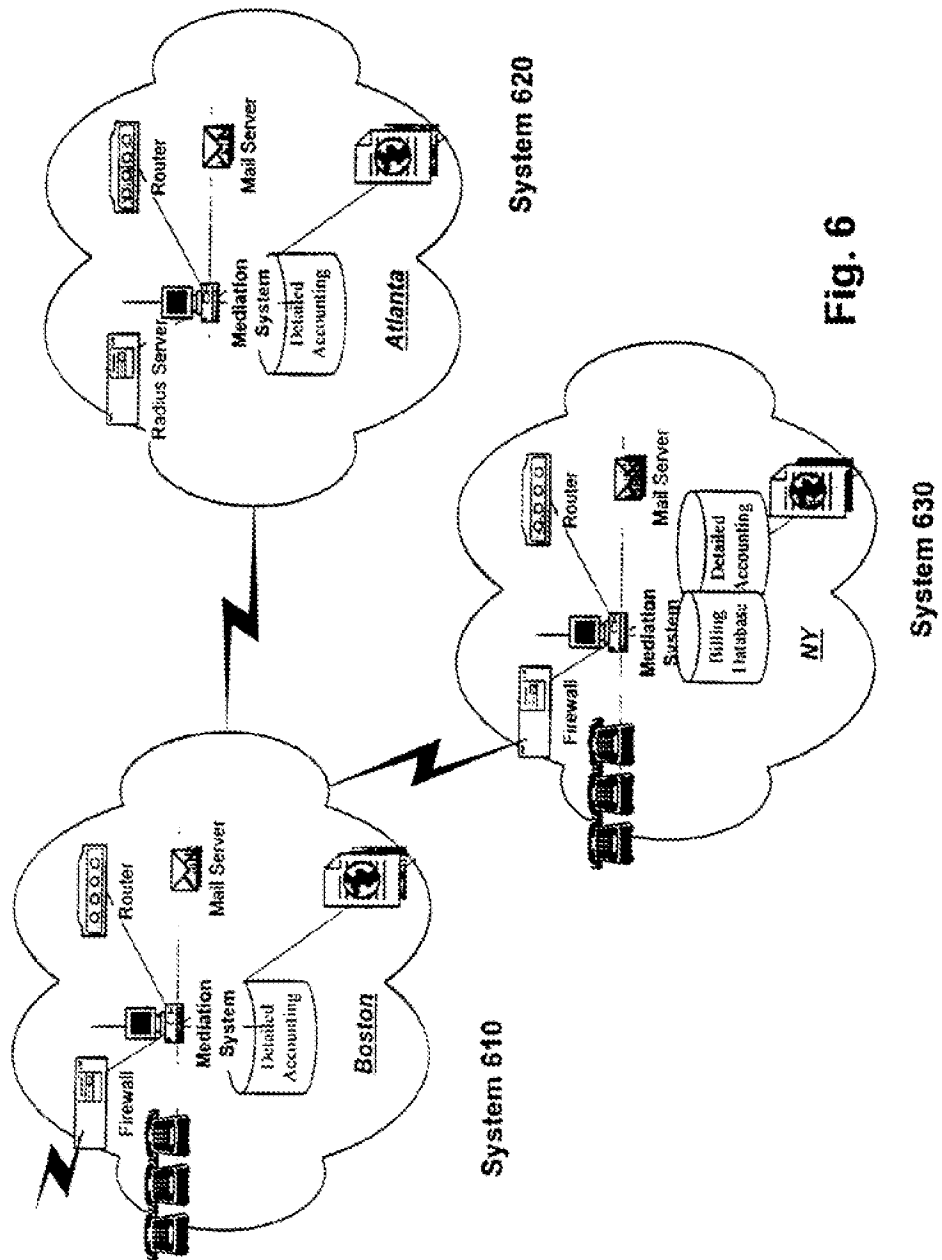


Fig. 6